The Standard Manual for Engineering Creature Comforts

Volume Two — The Internal Instrument

Michael L. Curzi · 36N9 GENETICS LLC

Contents

Volume Two — The Internal Instrument **6**

Part VI — The Essence Law **7**

Chapter 29 . The Essence Law **8**
The Law . . . 8
The Engineering . . . 10
The Specification . . . 21
The Failure Modes . . . 22
Worked Example . . . 24

Chapter 30 . The Interdisciplinary Bridge **28**
The Law . . . 28
The Engineering . . . 30
The Specification . . . 42
The Failure Modes . . . 43
Worked Example . . . 45

Chapter 31 . Sonic Alchemy in Built Space **49**
The Law . . . 49
The Engineering . . . 51
The Specification . . . 61
The Failure Modes . . . 62
Worked Example . . . 64

Chapter 32 . The Conversion Keys **71**
The Law . . . 71
The Engineering . . . 74
The Specification . . . 90
The Failure Modes . . . 91
Worked Example . . . 92

Chapter 33 . Harmonic Lattices and Essence Sets **96**
The Law . . . 96
The Engineering . . . 99

The Specification . . . 114
The Failure Modes . . . 115
Worked Example . . . 116

Part VII — The Perceptual Lens **119**

Chapter 34 . The Moment Equation **120**
The Law . . . 120
The Engineering . . . 127
The Specification . . . 140
The Failure Modes . . . 141
Worked Example . . . 144

Chapter 35 . The Calibration Ratio **148**
The Law . . . 148
The Engineering . . . 156
The Specification . . . 168
The Failure Modes . . . 170
Worked Example . . . 173

Chapter 36 . How Mathematical Functions Operate From a Perceptual Lens **178**
The Law . . . 178
The Engineering . . . 187
The Specification . . . 199
The Failure Modes . . . 200
Worked Example . . . 203

Chapter 37 . Subjective Reality as Postulate **207**
The Law . . . 207
The Engineering . . . 214
The Specification . . . 220
The Failure Modes . . . 221
Worked Example — THE MORNING RESERVOIR . . . 224

Chapter 38 . The Register Perception Actually Uses **229**
The Law . . . 229
The Engineering . . . 237
The Specification . . . 243
The Failure Modes . . . 244
Worked Example — MEASURING A SUBJECTIVE MAGNITUDE END TO END . . . 246

Chapter 39 . The Internal Surplus **249**
The Law . . . 249
The Engineering . . . 255
The Specification . . . 264
The Failure Modes . . . 265
Worked Example — ONE OPERATOR, SIX FRAMINGS . . . 267

Part VIII — The Personal Formulary **272**

Chapter 40 . The Self as Storage **273**
The Law . . . 273
The Engineering . . . 292
The Specification . . . 303
The Failure Modes . . . 304
Worked Example . . . 307

Chapter 41 . The Launch Budget **312**
The Law . . . 312
The Engineering . . . 323
The Specification . . . 335
The Failure Modes . . . 336
Worked Example . . . 338

Chapter 42 . The Gradient and the Desire Line **343**
The Law . . . 343
The Engineering . . . 356
The Specification . . . 370
The Failure Modes . . . 371
Worked Example . . . 374

Chapter 43 . The Echo and the Return Difference **379**
The Law . . . 379
The Engineering . . . 385
The Specification . . . 394
The Failure Modes . . . 395
Worked Example . . . 396

Chapter 44 . The Contraction and Arrival **400**
The Law . . . 400
The Engineering . . . 407
The Specification . . . 416
The Failure Modes . . . 417
Worked Example . . . 418

Chapter 45 . The Personal Formulary **422**
The Law . . . 422
The Engineering . . . 427
The Specification . . . 457
The Failure Modes . . . 458
Worked Example . . . 459

Part IX — The Interpersonal Formulary **464**

Chapter 46 . The Dyad **465**
The Law . . . 465
The Engineering . . . 467
The Specification . . . 481
The Failure Modes . . . 482

Worked Example — A Two-Person Design Practice 484

Chapter 47 . The Group Moment **489**
The Law 489
The Engineering 492
The Specification 504
The Failure Modes 506
Worked Example — A Regional Authority 507

Chapter 48 . Doctrine-to-Data **512**
The Law 512
The Engineering 514
The Specification 527
The Failure Modes 528
Worked Example — A Two-Site Agency 529

Chapter 49 . Trust, Obligation, and Reciprocity as Quantities **533**
The Law 533
The Engineering 535
The Specification 548
The Failure Modes 549
Worked Example — An Eight-Firm Regional Supplier Consortium 550

Chapter 50 . The Internal/External Correspondence **556**
The Law 556
The Engineering 557
The Specification 573
The Failure Modes 574
Worked Example — One Clinician and Her Unit, Measured With One Instrument 575

Chapter 51 . The Interpersonal Formulary **579**
The Law 579
The Engineering — the formulary 583
The Specification — using and extending the formulary 616
The Failure Modes — of formulary use 617
Worked Example — A Regional Volunteer Rescue Association, Diagnosed Through Eleven Entries 618

Part X — Validation **623**

Chapter 52 . The Formula Validation Protocol **624**
The Law 624
The Engineering 627
The Specification 634
The Failure Modes 635
Worked Example 638

Chapter 53 . Why Division Requires a Ratio Scale **647**
The Law 647

The Engineering . . . 649
The Specification . . . 658
The Failure Modes . . . 659
Worked Example . . . 662

Chapter 54 . Specific Objectivity **665**
The Law . . . 665
The Engineering . . . 668
The Specification . . . 678
The Failure Modes . . . 679
Worked Example . . . 682

Chapter 55 . Replication, Preregistration, and the Falsification Suite **686**
The Law . . . 686
The Engineering . . . 688
The Specification . . . 702
The Failure Modes . . . 704
Worked Example . . . 706

Chapter 56 . The Research Program **712**
The Law . . . 712
The Engineering . . . 714
The Specification . . . 734
The Failure Modes . . . 735
Worked Example . . . 737

Part XI — The Old Notation **740**

Chapter 57 . Correspondence Tables as Relational Graphs **741**
The Law . . . 741
The Engineering . . . 744
The Specification . . . 753
The Failure Modes . . . 755
Worked Example . . . 756

Chapter 58 . The Sigil as Nonlinear Map **760**
The Law . . . 760
The Engineering . . . 763
The Specification . . . 775
The Failure Modes . . . 776
Worked Example . . . 777

Chapter 59 . Memory Palace, Ritual, and Periodic Forcing **781**
The Law . . . 781
The Engineering . . . 784
The Specification . . . 797
The Failure Modes . . . 799
Worked Example . . . 801

Chapter 60 . This Is How It Is Done **806**
The Law 806
The Engineering 809
The Specification 821
The Failure Modes 822
Worked Example 824

Volume Two — The Internal Instrument

Part VI — The Essence Law

Chapter 29 . The Essence Law

Every chapter before this one has been about *quantity*: how much order is stored, how fast it leaks, how much difference a roster actually contains. This part of the manual is about *identity* — about giving a relationship a name that can be heard rather than a number that must be read.

The instrument is a single formula published in 2022 as *The Effective Method of Sonic Chemistry*. Its original application was the periodic table; its actual reach is wider, and this chapter establishes exactly how much wider and exactly how much less. We strip the formula to its algebraic core, verify that core against every number the source published, prove two structural theorems about what it can and cannot produce, and mark — in the open, with its badge on — the precise point at which mathematics stops and postulate begins.

The single most important sentence in this chapter is the one that limits it: the output of the essence law is an ESSENCE frequency, not a physical resonance. No material, molecule, element, room or person is claimed to physically resonate at any value computed here. The law assigns a pitch to a *position in an ordered category scheme*. What you do with that pitch is Chapter 31's business. What the pitch *is* — and is not — is this chapter's.

The Law

29.1 Statement

POSTULATE (The Essence Law). Let a discipline be given an explicit, ordered, finite set of essence types. Let n be the ordinal position — the **essence index** — of one such type. Then the essence frequency of that type is

```
f(n) = k * n^2                                          [Hz]

k = (9/8)^2 / phi^2 = (81/64) / (phi + 1) = (81/64) * (2 - phi)

phi = (1 + sqrt(5)) / 2 = 1.6180339887498948...

k  = 0.483425732988414 Hz   (0.48342573298841433274106978775598925726 to 40 places)
```

The identity `phi^2 = phi + 1` makes the second and third forms exact restatements of the first; `1/(phi+1) = 2 - phi = 0.3819660112501051`5, computed and confirmed below.

Two things in that block are proved and two are postulated, and they must never be confused.

Element	Status	Warrant
`f(n) = k n^2` is the exact collapse of the published recipe `((n/phi)*1.125)^2`	**PROVED** (algebra)	§29.2
`k = (81/64)/(phi+1) = 0.483425732988414`	**PROVED** (arithmetic)	§29.2, computed

Element	Status	Warrant
Every interval `f(m)/f(n)` is `m^2/n^2` and is independent of k	**PROVED** (Result One)	§29.4
A set indexed `1..N` has exactly `ceil(N/2)` pitch classes	**PROVED** (Result Two)	§29.5
That essence types *have* a frequency at all	**POSTULATE**	author's, 2022
That the frequency law is *square* in the index	**POSTULATE**	author's, 2022
That k is `(9/8)^2/phi^2` rather than any other constant	**POSTULATE** (design choice)	author's, 2022
That a given ordering of categories is the right one	**REQUIRES JUSTIFICATION**	Paper B §1.2

The mathematics in the middle two rows is unconditional: it is true of k n^2 for *any* positive k, and would remain true if the author had chosen k = 1. The mysticism is entirely in rows five through seven, and it is permitted here on the house condition established in Book One — **permitted in the open, with its badge on.**

29.2 The honesty clauses

These are mandatory and are restated in Chapters 30 and 31 wherever the law is used. They are not decoration. They are the warrant that lets this part of the manual sit inside a book whose other twenty-eight chapters are built on proved theorems.

H1 — Essence, not resonance. The output is an essence frequency. No claim is made, and none may be implied, that any material, molecule, element, institution, room or person physically resonates at these values. When Chapter 31 builds a room whose axial modes sit at essence frequencies, the room genuinely resonates — because a room of that dimension genuinely has that mode, by ordinary acoustics — but the *essence* does not, and the room's mode is a design target we chose, not a property of the essence we discovered.

H2 — k is chosen, not measured. Its form — a Pythagorean interval divided by the golden section — is a design decision by the author of the 2022 document. Nothing in the algebra requires it. Nothing in the intervals depends on it (§29.4). It sets the reference pitch and does no other work. Badge it as a postulate every time it appears.

H3 — The square law is a postulate about essence structure. That it coincides in form with the eigenvalue spacing of a one-dimensional infinite square well (E_n proportional to n^2) is a genuine and interesting mathematical correspondence, developed carefully in §29.8. **Coincidence of form is not evidence of mechanism, and this manual asserts no mechanism.**

H4 — The essence index requires a defensible ordering. Applying the law requires an explicit, justified ordinal on the discipline's categories, satisfying Paper B's validity conditions §1.2. An arbitrary ordering yields arbitrary frequencies. §29.6 gives the test.

H5 — Everything downstream inherits every clause. Chapter 30 layers the Interaction Surplus Framework on top of this law. That bridge inherits all of Papers A-E's honesty clauses, including that the surplus functional as deployed is a **lower bound** on an ensemble's information rather than the information itself, that Paper A is pure mathematics making no empirical or metaphysical claim, and that **no thermodynamic or physical consequence follows from any of it.**

The Engineering

29.3 The collapse

The 2022 source specifies a spreadsheet procedure. Column A holds the atomic number. A separate tab holds the golden ratio. Column C holds

```
(( AN / phi ) * 1.125)^2
```

with `AN` the atomic number and `1.125 = 9/8`. Expand it:

```
f(n) = ( (n / phi) * (9/8) )^2
     = ( n * (9/8) / phi )^2
     = n^2 * (9/8)^2 / phi^2
     = k * n^2                          with k = (9/8)^2 / phi^2
```

That is the whole derivation. There is nothing else in the recipe. **The published procedure is a pure square law with a constant of proportionality, and every appearance of the golden ratio and of the Pythagorean tone in the recipe is absorbed, once and for all, into that single constant.**

Numerically, computed to double precision and cross-checked at fifty decimal digits:

```
phi          = 1.618033988749895
phi^2        = 2.618033988749895
phi + 1      = 2.618033988749895        <- identical, as phi^2 = phi + 1
(9/8)^2      = 1.265625
1/(phi+1)    = 0.38196601125010515
2 - phi      = 0.3819660112501051

k = (9/8)^2 / phi^2   = 0.48342573298841435
k = (81/64)/(phi+1)   = 0.48342573298841435      difference exactly 0.0
k = (81/64)*(2-phi)   = 0.4834257329884143
k = (81/64)*(3-sqrt5)/2 = 0.4834257329884143

k to 50 digits = 0.48342573298841433274106978775598925726023369430832
```

All four closed forms agree. The house citation form is `k = 0.483425732988414 Hz`, fifteen significant figures, which is the precision at which the source document published.

29.4 Verification against the source document's own numbers

A collapse claim is worth nothing unverified. The source publishes 118 base frequencies — one per element of the periodic table, hydrogen through oganesson. Each was recomputed here from

`f(n) = k n^2` using exact fifty-digit arithmetic and compared with the published string.

Result: all 118 published values are reproduced by `k n^2`. The worst relative deviation across the whole table is 9.00e-15, at n = 17 (chlorine) — that is agreement to fourteen significant figures, which is the resolution limit of IEEE double-precision arithmetic.

n	element	published (2022 source)	`k * n^2` recomputed	relative deviation
1	H	0.483425732988417	0.483425732988414	5.5e-15
2	He	1.93370293195367	1.93370293195366	6.6e-15
3	Li	4.35083159689575	4.35083159689573	4.8e-15
8	O	30.9392469112587	30.9392469112585	5.9e-15
12	Mg	69.6133055503321	69.6133055503317	6.3e-15
16	S	123.756987645035	123.756987645034	7.5e-15
26	Fe	326.79579550017	326.795795500168	5.8e-15
29	Cu	406.561041443259	406.561041443256	6.3e-15
50	Sn	1208.56433247104	1208.56433247104	3.4e-15
64	Gd	1980.11180232056	1980.11180232055	7.5e-15
79	Au	3017.05999958071	3017.05999958069	5.4e-15
92	U	4091.71540401396	4091.71540401394	5.2e-15
100	Fm	4834.25732988417	4834.25732988414	5.5e-15
118	Og	6731.21990613072	6731.21990613068	5.8e-15

Note the structure that jumps out of the first and second-to-last rows: `f(100) = 10000 * f(1)` exactly, because `100^2 = 10^4`, so the two published strings carry the same digits with the decimal point moved four places. That is a signature of a pure power law and is visible by eye in the source document itself.

A forensic aside, because it changes what "verified" means. The published values are not merely close to `k n^2` — they are *systematically* high by about 2.8 parts in 10^15, too consistent to be rounding noise. Inverting the arithmetic to recover the `phi` the spreadsheet actually used:

```
phi implied by published f(1) = 1.61803398874989038...
true phi                      = 1.61803398874989484...
ratio                         = 1 - 2.76e-15
```

The implied value is exactly `phi` truncated to fifteen significant figures: `1.61803398874989`. Recomputing the entire table with that literal:

```
k(phi = 1.61803398874989) = 0.48342573298841723
-> 109 of 118 published values match to ALL FIFTEEN displayed digits
-> worst relative deviation 3.43e-15 (last-place float noise on the remaining nine)
```

So the collapse is not approximate. It is exact, and the residual discrepancy is fully explained by the spreadsheet having stored a fifteen-digit `phi` rather than a full-precision one. We can say that the published table *is* `k n^2`, and we can say precisely which `k`. This is the standard this manual holds itself to: not "the numbers look right" but "here is the arithmetic that produced them and here is the residue accounted for."

One further textual erratum, recorded because it will otherwise trip any reader who reimplements the method. The source document instructs: *"In one cell type in the equation 1+5^.5/2"*. Under standard operator precedence that expression evaluates to `1 + (sqrt(5)/2) = 2.118033988749895`, which is not `phi`. Recomputing `f(1)` with it gives `0.28212345562925867`, which matches nothing in the published table. **The spreadsheet plainly contained `(1+5^0.5)/2`; the document's prose dropped the parentheses.** Use `(1 + sqrt(5))/2`.

29.5 The constant, and what it is for

k is a Pythagorean interval divided by the golden section, and it is worth being precise about which. `9/8` is the **Pythagorean whole tone** — `(3/2)^2 / 2`, 203.910 cents. Its square, `81/64`, is the **Pythagorean ditone**, the major third of three-limit tuning, 407.820 cents, sharper than the just major third `5/4` (386.314 cents) by the syntonic comma. `phi^2 = phi + 1 = 2.618033988749895`. So k is the ditone reduced by `phi^2`.

The source document offers a rationale — that elements are standing waves on a phi spiral with eight per 360 degrees, hence the factor `9/8` to "establish that symmetry breaking in the spatial orientation of the golden spiral". **That rationale is a postulate. It is recorded because it is the author's stated reason, and badged because no derivation of it exists.**

What matters engineering-wise is the subject of the next section: *none of this survives into any relationship the law produces.* k sets the tuning fork. It sets nothing else.

29.6 RESULT ONE — every interval is a ratio of squares, and k cancels

Theorem (Result One). For any two essence indices `m, n >= 1`,

```
f(m)     k * m^2     m^2
---- =   -------  =  ---
f(n)     k * n^2     n^2
```

Proof. Immediate. k is a common factor of numerator and denominator. It cancels for every pair, for every `k > 0`. QED.

The proof is one line. The consequence is the whole chapter.

The golden ratio and the 9/8 set the absolute reference pitch and NOTHING ELSE. Every relationship in the system is a ratio of squared integers. If you replace k with 1, or with 440, or with the mass of the sun, every interval in the system is unchanged. `phi` does no relational work. It is a transposition, and a transposition is the one musical operation that changes nothing structural.

This is a demotion and it should be felt as one. A reader coming to this material from the wider literature of golden-ratio numerology will expect `phi` to be doing something. It is not. The harmonic content of the essence law is generated entirely by the integers.

Here is the interval lattice, computed. "Cents" is `1200*log2(ratio)`; "folded cents" is that value modulo 1200, i.e. the pitch class; "dev" is the deviation from the nearest twelve-tone equal-tempered interval.

m:n	ratio m^2/n^2	decimal	octaves	cents	folded cents	nearest 12-TET	dev (cents)
2:1	4/1	4.000000	2.0000	2400.00	0.00	unison	+0.00
3:2	9/4	2.250000	1.1699	1403.91	203.91	major 2nd	+3.91
4:3	16/9	1.777778	0.8301	996.09	996.09	minor 7th	-3.91
5:4	25/16	1.562500	0.6439	772.63	772.63	minor 6th	-27.37
6:5	36/25	1.440000	0.5261	631.28	631.28	tritone	+31.28
5:3	25/9	2.777778	1.4739	1768.72	568.72	tritone	-31.28
8:5	64/25	2.560000	1.3561	1627.37	427.37	major 3rd	+27.37
3:1	9/1	9.000000	3.1699	3803.91	203.91	major 2nd	+3.91
4:1	16/1	16.000000	4.0000	4800.00	0.00	unison	+0.00
5:2	25/4	6.250000	2.6439	3172.63	772.63	minor 6th	-27.37
7:4	49/16	3.062500	1.6147	1937.65	737.65	perfect 5th	+37.65
7:5	49/25	1.960000	0.9709	1165.02	1165.02	unison	-34.98
7:6	49/36	1.361111	0.4448	533.74	533.74	perfect 4th	+33.74
8:7	64/49	1.306122	0.3853	462.35	462.35	perfect 4th	-37.65
9:8	81/64	1.265625	0.3399	407.82	407.82	major 3rd	+7.82
16:9	256/81	3.160494	1.6601	1992.18	792.18	minor 6th	-7.82
10:9	100/81	1.234568	0.3040	364.81	364.81	major 3rd	-35.19
11:8	121/64	1.890625	0.9189	1102.64	1102.64	major 7th	+2.64
13:8	169/64	2.640625	1.4009	1681.06	481.06	perfect 4th	-18.94
15:8	225/64	3.515625	1.8138	2176.54	976.54	minor 7th	-23.46
5:1	25/1	25.000000	4.6439	5572.63	772.63	minor 6th	-27.37

Read the first four rows and the character of the system is visible. **The essence law doubles the distance of every classical interval.** A frequency ratio of `3:2` — the perfect fifth of ordinary acoustics — becomes `9:4`, an octave plus a major second, because squaring doubles the logarithm. The perfect fourth `4:3` becomes `16:9`, a minor seventh. The octave `2:1` becomes two octaves. This is a general fact, not a coincidence: since pitch interval is logarithmic and the law squares,

```
interval in cents between indices m and n = 2 * (classical interval between m and n)
```

Squaring the index doubles every interval in cents. The essence law is the classical harmonic series stretched by a factor of exactly two in log-frequency space. That single sentence predicts every row of the table above and is worth carrying.

The upper-triangular lattice for indices 1 through 6, each cell giving exact ratio, cents, and folded cents:

```
      n=1              n=2              n=3              n=4              n=5
m=2   4/1  2400.0    0.0
```

```
 m=3   9/1  3803.9  203.9   9/4  1403.9  203.9
 m=4  16/1  4800.0    0.0   4/1  2400.0    0.0  16/9   996.1  996.1
m=5  25/1  5572.6  772.6  25/4  3172.6  772.6  25/9  1768.7  568.7  25/16  772.6  772.6
m=6  36/1  6203.9  203.9   9/1  3803.9  203.9   4/1  2400.0   0.0   9/4 1403.9  203.9  36/25  6
```

Notice how often the folded-cents entry repeats: 0.0 at (2,1), (4,1), (4,2); 203.9 at (3,1), (3,2), (6,1), (6,2), (6,4). That repetition is not an accident of the display. It is Result Two.

29.7 RESULT TWO — octave folding is governed by the odd part of the index

Musical pitch is naturally taken modulo the octave: two tones an octave apart share a *pitch class* and are heard as "the same note, higher". The question "how many distinct notes does an essence set contain?" is therefore the question "how many distinct pitch classes does it contain?", and the essence law answers it with unusual sharpness.

Lemma. `f(2n) = k(2n)^2 = 4 k n^2 = 4 f(n)`. Doubling the index multiplies the frequency by four, which is exactly two octaves. Therefore `n` and `2n` always share a pitch class.

Theorem (Result Two). Write every index uniquely as `n = 2^a * m` with `m` odd (`m` is the *odd part* of `n`, `a` its 2-adic valuation). Then

```
f(n) = k * (2^a * m)^2 = k * 4^a * m^2 = f(m) * 4^a = f(m) * 2^(2a)
```

so `f(n)` is exactly `2a` octaves above `f(m)`. **Two essences share a pitch class if and only if their indices have the same odd part.**

Proof of "only if". Suppose `n = 2^a m` and `n' = 2^b m'` share a pitch class, i.e. `f(n)/f(n') = 2^j` for some integer `j`. Then `4^a m^2 / (4^b m'^2) = 2^j`, so `m^2/m'^2 = 2^(j − 2a + 2b)`. The left side is a ratio of odd squares, hence an odd rational in lowest terms; the right side is a power of two. The only odd power of two is `2^0 = 1`, so `m^2 = m'^2` and, since both are positive, `m = m'`. QED.

Corollary. An essence set indexed `1..N` contains exactly `ceil(N/2)` distinct pitch classes.

Proof. The odd part of any `n <= N` is odd and at most `n <= N`, so every realised class label is an odd number in `[1, N]`. Conversely every odd `m <= N` is itself in the index set and is its own odd part, so every such label is realised. Hence the class labels are exactly the odd integers in `[1, N]`, of which there are `ceil(N/2)`. QED.

Here is the full index-to-class table for `N` up to 24, computed:

n	n = 2^a * m	odd part m	a	f(n) = k n^2 (Hz)	class representative f(m) (Hz)
1	2^0 * 1	1	0	0.483426	0.483426
2	2^1 * 1	1	1	1.933703	0.483426
3	2^0 * 3	3	0	4.350832	4.350832
4	2^2 * 1	1	2	7.734812	0.483426
5	2^0 * 5	5	0	12.085643	12.085643
6	2^1 * 3	3	1	17.403326	4.350832
7	2^0 * 7	7	0	23.687861	23.687861
8	2^3 * 1	1	3	30.939247	0.483426

n	n = 2^a * m	odd part m	a	f(n) = k n^2 (Hz)	class representative f(m) (Hz)
9	2^0 * 9	9	0	39.157484	39.157484
10	2^1 * 5	5	1	48.342573	12.085643
11	2^0 * 11	11	0	58.494514	58.494514
12	2^2 * 3	3	2	69.613306	4.350832
13	2^0 * 13	13	0	81.698949	81.698949
14	2^1 * 7	7	1	94.751444	23.687861
15	2^0 * 15	15	0	108.770790	108.770790
16	2^4 * 1	1	4	123.756988	0.483426
17	2^0 * 17	17	0	139.710037	139.710037
18	2^1 * 9	9	1	156.629937	39.157484
19	2^0 * 19	19	0	174.516690	174.516690
20	2^2 * 5	5	2	193.370293	12.085643
21	2^0 * 21	21	0	213.190748	213.190748
22	2^1 * 11	11	1	233.978055	58.494514
23	2^0 * 23	23	0	255.732213	255.732213
24	2^3 * 3	3	3	278.453222	4.350832

And the class count, computed and checked against `ceil(N/2)` for every `N` up to 24:

N	distinct odd parts realised	count	ceil(N/2)
1	1	1	1
2	1	1	1
3	1,3	2	2
4	1,3	2	2
5	1,3,5	3	3
6	1,3,5	3	3
7	1,3,5,7	4	4
8	1,3,5,7	4	4
9	1,3,5,7,9	5	5
10	1,3,5,7,9	5	5
11	1,3,5,7,9,11	6	6
12	1,3,5,7,9,11	6	6
13, 14	1,3,5,7,9,11,13	7	7
15, 16	1,3,5,7,9,11,13,15	8	8
17, 18	1,3,…,15,17	9	9
19, 20	1,3,…,17,19	10	10
21, 22	1,3,…,19,21	11	11
23, 24	1,3,…,21,23	12	12

Twelve essences yield only six pitch classes. That is the available harmonic vocabulary, and it is a hard structural limit, not a stylistic choice. The compression ratio is exactly 2 for every `N`, which is the flattest possible statement of the constraint: *half of any essence set is harmonically*

invisible. You may have 118 elements; you have 59 notes. You may have 24 departments; you have 12 notes. Doubling the roster does not double the palette in any way you can hear beyond the doubling already accounted for.

29.8 The pitch-class lattice is generated by the odd primes

Result Two admits a sharper algebraic statement that is useful when you are designing a set rather than analysing one.

Octave-reduce each odd class label `m` by dividing `m^2` by the largest power of two below it. The result lies in `[1,2)`:

odd m	m^2	a	folded ratio m^2/2^a	cents	nearest 12-TET	dev
1	1	0	1.000000000	0.00	unison	+0.00
3	9	3	1.125000000	203.91	major 2nd	+3.91
5	25	4	1.562500000	772.63	minor 6th	-27.37
7	49	5	1.531250000	737.65	perfect 5th	+37.65
9	81	6	1.265625000	407.82	major 3rd	+7.82
11	121	6	1.890625000	1102.64	major 7th	+2.64
13	169	7	1.320312500	481.06	perfect 4th	-18.94
15	225	7	1.757812500	976.54	minor 7th	-23.46
17	289	8	1.128906250	209.91	major 2nd	+9.91
19	361	8	1.410156250	595.03	tritone	-4.97
21	441	8	1.722656250	941.56	major 6th	+41.56
23	529	9	1.033203125	56.55	minor 2nd	-43.45

Because `m` is odd, factor it as `m = p1^e1 * p2^e2 * ...` over odd primes. Then `m^2 = p1^(2e1) * p2^(2e2) * ...`, and the folded class of `m` is the product of the folded classes of `p_i^2` taken `e_i` times, modulo octaves. Therefore:

The essence pitch-class lattice is the free abelian group generated by the octave-reductions of the squares of the odd primes. The generators, in ascending prime order:

```
prime 3  -> 9/64 folded to 9/8    = 203.910 cents    (Pythagorean whole tone)
 prime 5   ->  25/16                  = 772.627 cents
 prime 7   ->  49/32                  = 737.652 cents
 prime 11  ->  121/64                 = 1102.636 cents
 prime 13  ->  169/128                = 481.062 cents
```

Check the composite classes against the generators: class 9 = 3^2, so its folded class must be `(9/8)^2 = 81/64 = 1.265625`, which is exactly the table's entry and `2 * 203.910 = 407.820` cents. Class 15 = 3*5, so (9/8)*(25/16) = 225/128, folded 225/128 = 1.7578125, and 203.910 + 772.627 = 976.537 cents, matching the table's 976.54. Class 21 = 3*7*: `203.910 + 737.652 = 941.562`, matching 941.56.

This is not decoration. It has a direct design consequence: **an essence set that runs `1, 2, 4, 8, 16, ...` produces exactly one note; an essence set whose indices are the odd primes produces the maximally independent harmonic vocabulary available at that size.** Where an

ordinal scheme admits legitimate freedom (§29.9 says how much), prime indices buy harmonic independence and powers of two throw it away.

Note also a genuine curiosity: 9/8 appears twice here for unrelated reasons. It is inside k because the author put it there, and it is the folded class of index 3 because 3^2 = 9 and 2^3 = 8. The second survives any change of k; the first does not.

Four exact design laws follow once the lattice has to be rendered, and they are constraints rather than preferences. They are derived in Chapter 71 from the Gabor time-bandwidth limit and from the arithmetic above, and they are stated here because they bound what an essence set may be before anyone chooses one. All four are computed, not asserted.

Law 1 — the audible-set ceiling. With a single integer octave transposition f ∗ 2^j and a band of 20 Hz to 20 kHz, the whole set n = 1..N fits only if k ∗ 2^j >= 20 and k ∗ N^2 ∗ 2^j <= 20000. The first forces j >= 6; the second then caps N. **At most N = 25 essences fit under a single transposition, and only at j = 6** — j = 7 admits 17, j = 8 admits 12. **A twelve-house set fits comfortably. A thirty-two-system set does not, and requires split transposition**, which is a real design decision with real consequences and must be declared wherever it is used.

Law 2 — the minimum-gap cliff. Fold the set's pitch classes into one octave and measure the smallest gap. For N from 7 to 16 it is exactly **50/49, or 34.976 cents** — the near-coincidence of classes 7 and 5 in the table above. **The moment index 17 enters the set, the minimum gap collapses to exactly 289/288, or 6.001 cents**, because 17^2 = 289 lands beside 9/8 = 288/256, the folded class of index 3. **The boundary is exact and it is arithmetic: a set of sixteen and a set of seventeen are different design regimes.**

Three ratios, and they are not the same number. The cliff is quoted three ways in this manual and each is correct in its own units, so state which one you mean:

Quantity	Value	What it measures
Gap in **cents**	ln(50/49)/ln(289/288) = **5.828**	how much smaller the interval is, logarithmically
Cycles of the tone	288/49 = **5.878**	3.899 cycles becomes 22.918 (Law 3)
Achievable **rate**	**8, exactly**	symbols per second, and this is the design limit

The rate figure is the one that binds, and it is exactly 8 rather than either of the others. The reason is that the two binding pairs do not sit at the same height in the octave: the 50/49 pair straddles 49/32, the 289/288 pair straddles 9/8 = 288/256. Gabor resolution depends on the **absolute** gap in hertz, which is f/32 before the cliff and f/256 after — a ratio of exactly 8 at every base pitch. Dividing fractional gaps instead (288/49) silently assumes both pairs sit at the same frequency, and they do not. See Chapter 71 §71.12, which derives this.

Law 3 — the surplus-resolution bound. Resolving a difference dF in surplus requires at least **1 / (8 pi dF) cycles** of the rendered tone, and the bound is independent of the transposition because it constrains a ratio. At dF = 0.01 nats that is 3.98 cycles; at dF = 0.001, 39.79. **This is the exact Gabor bridge to the Surplus Law: it prices, in time, what it costs to hear a difference in surplus.** Chapter 30 §30.9 works the consequence for measurement precision.

Law 4 — the rectangular gate is disqualified. A rectangular on/off envelope has **no finite spectral second moment** — its `sinc^2` tails fall as `1/f^2` and the second-moment integral is numerically divergent — so it has no defined bandwidth and cannot serve as a rendering envelope at all. **Any implementation must use a windowed envelope.** Chapter 31 specifies one.

29.9 What an essence index is, and what makes an ordering defensible

An **essence index** is the ordinal position of an essence type within its discipline's ordered category set. Atomic number is one instance of an essence index. It is not the definition of one, and the 2022 source's decision to apply the formula to chemistry is an application, not a restriction.

Atomic number is, however, the *best possible* case, and understanding why tells you what you need in every other case. It is **intrinsic** (a count of protons, a property of the thing rather than of our filing system), **total** (every element has exactly one and no two share one), **dense from 1** (no gaps, so `ceil(N/2)` applies without qualification), **independently re-derivable** (anyone can recount and get the same integer), and **stable** (it does not change when the discipline reorganises its textbooks).

Most disciplines have nothing like this. Paper B, *Modeling Interpretation*, states the conditions under which the surplus framework's formal structure may be applied to a target domain at all, and those conditions govern here too. Paper B §1.2 requires, before any modeling map is applied:

1. **State encoding** — a map from target-domain items to normalised vectors, so every represented state has a determinate location in the model.
2. **Category assignment** — a finite partition into idealised categories, represented by blocks `H_1, ..., H_N`.
3. **Geometric dependence** (the S1 analogue) — interaction modelled through a scalar similarity parameter such as an angle, cosine, or inner-product-derived quantity.
4. **Zero baseline** (S2 analogue) — parallel or representationally equivalent states correspond to zero modelled surplus.
5. **Affine effective-count law** (S3 analogue) — adopted as a modeling assumption, **not read directly off the world.**
6. **Normalisation** (S4 analogue) — the maximum modelled surplus is finite and depends on the chosen category count `N`.

Paper B is explicit about the exclusion that follows: a target domain lacking a stable encoding, a defensible category scheme, or an interpretable angle-like geometry **is outside the scope of the model.** The same exclusion applies to the essence law, and this manual enforces it.

To those six, the house method adds one more, specific to the essence law because the essence law needs an *order* and not merely a *partition*:

7. **Declared ordinal rule.** The rule that assigns index n to category `i` must be published alongside the frequencies, must be applicable by a third party without consulting you, and must include its tie-break procedure. If the ordering came from a list you happened to write down in that sequence, say so; that is a weak ordering and its weakness must travel with the numbers.

An ordering test you can run in ten minutes. Take your essence set. Permute two indices. Recompute. Report the change. Worked, for a twelve-set:

- Swap indices 3 and 12. Frequencies move from 4.350832 Hz to 69.613306 Hz and back — a factor of exactly 16, four octaves. **But the pitch classes do not change at all**, because 3 and 12 share the odd part 3. This permutation is harmonically invisible.
- Swap indices 3 and 5. Class 3 (folded ratio 9/8, 203.91 cents) becomes class 5 (25/16, 772.63 cents). The pitch class moves by `772.63 − 203.91 = 568.72` cents — nearly half an octave, a tritone's worth of change, from swapping two adjacent list positions.

If a plausible re-ordering of your categories moves a pitch class by half an octave, your ordering is load-bearing and must be justified before any of the frequencies are published. That is the test. It costs nothing and it is the single most common thing practitioners skip.

29.10 The composition rule, and its collapse

The 2022 source also gives a rule for combining essences. For H20: three atoms total, so raise the hydrogen frequency to the 2/3 power, the oxygen frequency to the 1/3 power, multiply, and square:

```
f(H2O) = ( f(H)^(2/3) ∗ f(O)^(1/3) )^2
```

This too collapses. Let a compound contain essence indices `n_i` with multiplicities `c_i`, and let `T = sum(c_i)` be the total count. Define `G` as the count-weighted geometric mean of the indices:

```
G = ( prod_i  n_i^(c_i) )^(1/T)
```

Then

```
f_compound = ( prod_i f(n_i)^(c_i/T) )^2
           = ( prod_i (k n_i^2)^(c_i/T) )^2
           = ( k^(sum c_i / T) ∗ prod_i n_i^(2 c_i / T) )^2
           = ( k ∗ G^2 )^2
           = k^2 ∗ G^4
```

The composition rule is k^2 times the fourth power of the count-weighted geometric mean of the constituent essence indices. Test it on water: `G = (1^2 ∗ 8^1)^(1/3) = 8^(1/3) = 2.00000000`, so `f = k^2 ∗ 16 = 3.7392070290462`, against the source's published `3.73920702904621`. Agreement to fourteen significant figures.

Verified across every composite value the source publishes — 24 pharmacological compounds, 8 nucleotides in the carbon column, and 6 each in the silicon- and germanium-substituted columns, 44 values in all:

37 of 44 published composite values are reproduced by `f = k^2 G^4` to better than one part in 10^12. Seven disagree. The seven, with the size of the disagreement:

formula	published	k^2 G^4	relative deviation
O2	11.1246117974981	957.23699944	85.05
C5H6N2O2 (thymine)	4.66961137812091	21.80527042	3.670
C12H17N2O4P	18.5343343733681	13.36731640	0.279
C257H383N65O77S6	14.2762378831273	11.27999043	0.210
C43H66N12O12S2	12.8278625485662	11.29835161	0.119
C77H120N18O26S	11.0756447985031	10.43459440	0.058

formula	published	k^2 G^4	relative deviation
C158H251N39O46S	9.46036014007241	9.41751106	0.005

These seven are spreadsheet errors in the source, not evidence of a different rule, and the O2 row proves it. For a single-element species the count-weighted geometric mean is the element index itself, independent of how many atoms are present: `G = 8` for 0, 02, 03, 0_x, always. Therefore `f = k^2 * 8^4 = 957.23699944` for *any* pure-oxygen formula, and **no atom count whatever can produce the published 11.1246117974981.** That row cannot have come from the stated rule. A rule error would fail uniformly across the table; a data-entry error fails on scattered rows. We observe scattered rows. The rule is confirmed and seven cells are errata.

Two engineering notes, both honest flags:

The squaring destroys dimensional consistency. `f(n)` is quoted in Hz; `k^2 G^4` carries Hz-squared if the units are taken seriously. The source reads the squaring as a change of *structural level* — "each time you square the result you establish a level of harmonic molecular structure" — which is a reasonable reading of a designed formalism and is what this manual adopts. **But the composite number's "Hz" label is a convention, not a dimensional quantity, and it must not enter an acoustic calculation without an explicit re-anchoring step.** Chapter 31 re-anchors by octave folding, a ratio operation and therefore unit-agnostic; that is the safe path.

Composition flattens the vocabulary further. Because `G` is a geometric mean, low-index constituents pull it toward 1 and high-index constituents raise it slowly. Across the 24 pharmacological compounds in the source, `G` ranges only from 1.84 to 4.30 — a factor of 2.34, which is a factor of 30 in `k^2 G^4`, or about 4.9 octaves. **All twenty-four compounds occupy under five octaves of essence pitch and eighteen of them occupy under two.** Composition is strongly contractive on the pitch space and should be expected to cluster, not disperse.

29.11 The infinite square well: a correspondence of form

A particle of mass `m` confined to a one-dimensional box of width `L` with infinitely high walls has energy eigenvalues

```
E_n = n^2 h^2 / (8 m L^2)          n = 1, 2, 3, ...
```

Divide by Planck's constant to get the associated frequency:

```
nu_n = E_n / h = n^2 * h / (8 m L^2)
```

which is `constant * n^2`. **This is the same functional form as the essence law, with the constant `h/(8 m L^2)` playing the role of k.**

The correspondence can be pushed to a specific number. Setting `h/(8 m_e L^2) = k` with `m_e` the electron mass and CODATA values `h = 6.62607015e-34 J s` (exact, by SI definition) and `m_e = 9.1093837015e-31 kg`:

```
L = sqrt( h / (8 * m_e * k) ) = 0.013714300 m = 1.3714 cm

check: h / (8 * m_e * L^2) = 0.483425732988414    <- equals k
```

So there exists a box, 1.37 centimetres wide, for which an electron's level frequencies coincide numerically with the essence frequencies of a set indexed from 1.

This is a coincidence of form and nothing more, and the manual says so in the same breath in which it states the correspondence. Three reasons:

1. **k was chosen, not measured.** Because k is a free constant, a box width can always be found to match it. The existence of L is not a discovery; it is a division. Had k been `(4/3)^2/phi^3` the box would be a different size and the "correspondence" equally available. **A correspondence that can be manufactured for any constant carries no information about any constant.**
2. **The 1.37 cm box has no physical significance.** Not the size of an atom (five orders of magnitude too big), not a lattice spacing, not a wavelength of anything in the problem.
3. **`n^2` is the most common eigenvalue pattern in elementary physics**, appearing wherever a wave is confined between fixed boundaries with quantised half-wavelengths. Sharing it is not remarkable.

What *would* be evidence of mechanism, stated plainly: a preregistered prediction, derived from the essence law before measurement, of a physical observable that the law gets right and no simpler model does, on a held-out sample, with the ordering rule fixed in advance. Chapter 25's protocol applies unmodified. **This manual makes no such prediction and offers no such evidence. The essence law is used here as a naming and design instrument, and nothing in Chapters 29 through 31 depends on it being anything more.**

The correspondence is still worth carrying for one legitimate reason: it tells you what the law's *shape* implies. In a square well, level spacing grows without bound while the *log*-spacing collapses; the low levels are where the interesting structure is. The essence law inherits exactly that character. The interval between neighbouring indices is `2 * 1200 * log2(n/(n−1))` cents: 1403.9 at `n = 2`, 177.7 at `n = 20`, 69.8 at `n = 50`. **The first few essence indices carry almost all of the harmonic interest; past about index 20 the set thins through equal temperament and out into microtones.** Design accordingly.

The Specification

Numbered, checkable requirements. A published essence-frequency set that fails any of these is not conformant to this manual.

ES-1 — State the law in its collapsed form. Publish `f(n) = k n^2` with `k = 0.483425732988414` Hz, not the spreadsheet recipe. The recipe hides that a single constant does all the reference-setting work, and it invites the parenthesis error documented in §29.4.

ES-2 — Badge k at every appearance. k is a chosen reference constant; its form is a design decision by the author of the 2022 source. Never present it as measured, derived, or fundamental.

ES-3 — Publish the ordinal rule with the frequencies. State how index n was assigned to category `i`, in terms a third party can apply without consulting you, including tie-breaks. A frequency table without its ordinal rule is not a result; it is a list.

ES-4 — Run and report the permutation sensitivity test. Swap at least three plausible index pairs, recompute, report the maximum pitch-class displacement in cents. If any displacement exceeds 100 cents, the ordering is load-bearing; say so and justify it. (§29.9)

ES-5 — Verify Paper B §1.2 conditions 1, 2 and 6 before applying the law. State the encoding, the partition, and `N`. If any of the three cannot be stated, the domain is out of scope.

ES-6 — Report `ceil(N/2)` alongside `N`. This prevents the commonest design error: planning a composition around `N` voices when only `ceil(N/2)` are distinguishable. (§29.7)

ES-7 — Never publish a frequency without its octave. Every value carries its raw Hz or its transposition exponent `j` with `f_published = f_raw * 2^j`, plus the band. Two values quoted without their `j` cannot be compared.

ES-8 — Quote to no more than fifteen significant figures. That is the source's precision and the limit of double-precision arithmetic. For any built application four figures already exceed the tuning resolution of every physical instrument you will use.

ES-9 — Restate honesty clause H1 in any document carrying a frequency table. "This is an essence frequency, not a physical resonance. No claim is made that any material, molecule, element, institution or person physically resonates at these values." On the same page as the numbers.

ES-10 — When composing, publish `G` and `T` with the result. `f = k^2 G^4` is the closed form; give the count-weighted geometric mean and the member count so the value is reproducible.

ES-11 — Re-anchor composite values by ratio, never by unit. Composite outputs carry `Hz^2`. Bring them into band by octave folding — a ratio operation — and state the fold. Never feed a composite value into an acoustic formula expecting Hz. (§29.10)

ES-12 — Note the square-well disclaimer in the same paragraph as the correspondence. Never in a later section, never in a footnote. Coincidence of form is not evidence of mechanism. (§29.11)

ES-13 — Recompute the source's published values before relying on them. Seven of the 44 composite values in the 2022 document are errata (§29.10). Recompute; do not transcribe.

ES-14 — Prefer prime indices when index assignment is genuinely free. Odd-prime indices maximise harmonic independence; powers of two collapse onto one class. Where the ordinal rule does *not* admit freedom, ES-3 governs and this requirement is void. (§29.8)

The Failure Modes

Each entry gives the fault, the parameter or claim it damages, its diagnostic signature, and the fix.

FM-29.1 THE ARBITRARY ORDINAL

Damages: every frequency in the set. *Description.* An ordering chosen for convenience — alphabetical, historical, the order items came up in a meeting — is fed into `f(n) = k n^2` and the output treated as intrinsic. *Signature.* Asked "why is Ecology index 10?", the answer is "because that's where it is on the sheet". *Fix.* ES-3 and ES-4. Publish the rule; run the permutation test.

If displacement is large and the rule is weak, **the correct output is a range, not a frequency.** Report the range.

FM-29.2 THE PHYSICAL RESONANCE CLAIM

Damages: the credibility of everything else in the volume. *Description.* "Iron resonates at 326.80 Hz." It does not. Iron's essence index is 26 and `k * 676 = 326.795795500168` — a number attached to a position in an ordering. *Signature.* Any sentence in which a material, molecule, body part or person is the *subject* of the verb "resonates" with an essence frequency as its object. *Fix.* H1, restated per ES-9. Rewrite so the essence set is the subject: "index 26 of the periodic ordering carries essence pitch class 13."

FM-29.3 THE UNBADGED CONSTANT

Damages: the reader's ability to tell proof from postulate. *Description.* `k = (9/8)^2/phi^2` is presented as though the golden ratio had been discovered in the data rather than placed in the formula by hand. *Signature.* Prose in which `phi` is described as "emerging", "appearing", or "governing". *Fix.* ES-2, plus Result One: state in the same passage that `k` cancels from every relationship and therefore governs nothing but the reference pitch.

FM-29.4 THE PHANTOM VOCABULARY

Damages: every downstream decision that assumed `N` distinguishable tones. *Description.* A twelve-category scheme is sonified as twelve voices; six are octave duplicates and the "twelve-part harmony" is a six-note chord doubled. *Signature.* Two "different" voices a listener cannot tell apart; two indices whose ratio is a power of four. *Fix.* ES-6. Design to `ceil(N/2)`. Twelve distinct tones require at least twenty-three justified essence indices.

FM-29.5 THE UNSTATED OCTAVE

Damages: comparability between any two published tables. *Description.* One table publishes raw values, another the same set folded into 220-440 Hz; compared directly they appear to disagree by orders of magnitude. *Signature.* Values from two sources differing by a clean power of two. *Fix.* ES-7. Every value carries its `j` and its band, always.

FM-29.6 PRECISION THEATRE

Damages: trust, and nothing else — but that is enough. *Description.* Publishing `326.795795500168` Hz as a tuning target for a physical object whose own frequency drifts a percent with temperature, and whose best tuner resolves 0.1 cent. *Signature.* More than about six significant figures on any value that will be built rather than computed. *Fix.* ES-8. Fifteen figures in the derivation table, four in the build specification.

FM-29.7 THE SQUARE-WELL SLIDE

Damages: the boundary between this chapter and physics. *Description.* The `n^2` correspondence is noted, then two pages later the essence law is called "quantum", and by the end someone is claiming energy levels. *Signature.* "Quantum", "eigenvalue" or "energy level" used without a

disclaimer in the same paragraph. *Fix.* ES-12 — and the harder fix: if the argument cannot be made without the physics vocabulary, it was leaning on borrowed authority and must be rewritten without it.

FM-29.8 COMPOSITION UNIT DRIFT

Damages: any acoustic calculation downstream. *Description.* A composite value from `k^2 G^4` — dimensionally Hz-squared — is fed into `L = c/(2f)` or a Sabine calculation as though it were Hz. *Signature.* A tube length or absorption target absurd by three or more orders of magnitude. *Fix.* ES-11. Fold into the target band first, then use ratios only.

FM-29.9 THE SPREADSHEET ORACLE

Damages: seven values in 44, unpredictably. *Description.* Transcribing published values instead of recomputing them, inheriting the source's data-entry errors, then defending them because they are "published". *Signature.* A value that resists reproduction by the stated rule while its neighbours reproduce to fourteen figures. *Fix.* ES-13. Where recomputation disagrees by more than 1e-12 relative, the published number is wrong until proven otherwise — and say so in print, as §29.10 does.

FM-29.10 THE POWERS-OF-TWO SET

Damages: the entire harmonic output. *Description.* An essence set built on indices `1, 2, 4, 8,` 16 — a natural doubling scheme. Every member has odd part 1. **The set has exactly one pitch class.** *Signature.* All voices in unison across octaves; a "chord" with no interval in it. *Fix.* ES-14 where index assignment is free; where it is not, accept and report the result — that a category scheme is harmonically degenerate is a real finding about the scheme.

Worked Example

Part 1 — The periodic table, first twenty elements, carried end to end

This is the source document's own domain and therefore the strongest available test: we can check every intermediate against a published value.

Step 1 — Establish the ordinal. The essence index is atomic number: the count of protons in the nucleus. It satisfies all five properties of §29.9 — intrinsic, total, dense from 1, independently re-derivable, stable. This is the best case and it is why the source chose it. Paper B §1.2 condition 2 is satisfied by the partition of matter into elements; condition 6 is satisfied with `N = 118`.

Step 2 — Compute raw essence frequencies. `f(n) = 0.`483425732988414 `∗ n^2`.

Step 3 — Compute odd parts and pitch classes. By Result Two.

Step 4 — Fold into an audible band. Band chosen: `[220, 440)` Hz — one octave, anchored on A3 = 220 Hz. `f_folded = f ∗ 2^j`, `j` the integer bringing `f` into the band. The band and `j` are published for every value, per ES-7.

n	element	f(n) (Hz)	odd part	class	j	f folded to [220,440)
1	H	0.483426	1	PC1	+9	247.5140
2	He	1.933703	1	PC1	+7	247.5140
3	Li	4.350832	3	PC3	+6	278.4532
4	Be	7.734812	1	PC1	+5	247.5140
5	B	12.085643	5	PC5	+5	386.7406
6	C	17.403326	3	PC3	+4	278.4532
7	N	23.687861	7	PC7	+4	379.0058
8	O	30.939247	1	PC1	+3	247.5140
9	F	39.157484	9	PC9	+3	313.2599
10	Ne	48.342573	5	PC5	+3	386.7406
11	Na	58.494514	11	PC11	+2	233.9781
12	Mg	69.613306	3	PC3	+2	278.4532
13	Al	81.698949	13	PC13	+2	326.7958
14	Si	94.751444	7	PC7	+2	379.0058
15	P	108.770790	15	PC15	+2	435.0832
16	S	123.756988	1	PC1	+1	247.5140
17	Cl	139.710037	17	PC17	+1	279.4201
18	Ar	156.629937	9	PC9	+1	313.2599
19	K	174.516690	19	PC19	+1	349.0334
20	Ca	193.370293	5	PC5	+1	386.7406

Step 5 — Count the vocabulary. `N = 20`, so `ceil(20/2) = 10` distinct pitch classes: `{1, 3, 5, 7, 9, 11, 13, 15, 17, 19}`. Ten notes from twenty elements. Confirmed against the class column: PC1 five times (H, He, Be, O, S), PC3 three (Li, C, Mg), PC5 three (B, Ne, Ca), PC7 twice (N, Si), PC9 twice (F, Ar), and PC11, PC13, PC15, PC17, PC19 once each — twenty in total.

Step 6 — Read the structure. Hydrogen, helium, beryllium, oxygen and sulphur are *the same note*. Not a metaphor: 1, 2, 4, 8 and 16 share the odd part 1, so their essence frequencies stand in ratios `1 : 4 : 16 : 64 : 256`, which is `0, 2, 4, 6, 8` octaves. **A tuning built on this set has one voice at PC1 sounding in five octaves and single voices at nine other classes — a strong fundamental with a thin upper structure, and that is a structural fact about the ordering, available before a single note is played.**

Step 7 — Verify and clause. Every raw frequency above was recomputed and matches the 2022 document to a worst relative deviation of 9.0e-15 across the full 118-element set (§29.4); nothing is transcribed. Per ES-9: *these are essence frequencies, not physical resonances; no claim is made that hydrogen, iron or any other element physically resonates at any value in this table.*

Part 2 — The same arithmetic on a non-chemical ordinal

The point of Chapter 29 is that atomic number is *an instance* of an essence index. Here is a second instance from this book's own lineage: **Book One's twelve nodes.** The node sequence is a designed process order — each node is a stage an operator passes through, and the order is the order of passage. That makes it a defensible ordinal under §29.9: intrinsic to the method, total, dense from 1, stable, and independently re-derivable by anyone holding the volume.

The four nodes this manual names elsewhere sit at indices 2 (GRADIENT), 7 (COLLAPSE), 9 (RECURSE) and 10 (FIXEDPOINT), with 11 (SEAL) — the remainder are referred to by index only, because inventing names for them would violate CANON §5.3.

`N = 12`, so `ceil(12/2) = 6` pitch classes: `{1, 3, 5, 7, 9, 11}`.

n	node	f(n) (Hz)	odd part	class	j	folded to [220,440)
1	node 01	0.483426	1	PC1	+9	247.5140
2	node 02 GRADIENT	1.933703	1	PC1	+7	247.5140
3	node 03	4.350832	3	PC3	+6	278.4532
4	node 04	7.734812	1	PC1	+5	247.5140
5	node 05	12.085643	5	PC5	+5	386.7406
6	node 06	17.403326	3	PC3	+4	278.4532
7	node 07 COLLAPSE	23.687861	7	PC7	+4	379.0058
8	node 08	30.939247	1	PC1	+3	247.5140
9	node 09 RECURSE	39.157484	9	PC9	+3	313.2599
10	node 10 FIXEDPOINT	48.342573	5	PC5	+3	386.7406
11	node 11 SEAL	58.494514	11	PC11	+2	233.9781
12	node 12	69.613306	3	PC3	+2	278.4532

Sorted by pitch and reduced to a single octave, the six classes form this scale:

#	class	folded ratio	cents above PC1	step to next (cents)
1	PC1	1.000000000	0.00	203.91
2	PC3	1.125000000	203.91	203.91
3	PC9	1.265625000	407.82	329.83
4	PC7	1.531250000	737.65	34.98
5	PC5	1.562500000	772.63	330.01
6	PC11	1.890625000	1102.64	97.36

The six steps sum to `203.91 + 203.91 + 329.83 + 34.98 + 330.01 + 97.36 = 1200.00` cents. The scale closes at the octave, as it must.

Three structural observations, all of them consequences of §29.8 and none of them stylistic:

1. **The scale opens with two exact Pythagorean whole tones.** PC1 -> PC3 -> PC9 is `1 : 9/8 : 81/64`, the tone-chain generated by the odd prime 3. Nodes 03, 06 and 12 (class 3) and node 09 (class 9) are related to node 01 by pure three-limit intervals.
2. **PC7 and PC5 are 34.98 cents apart — a third of a semitone.** Node 07 COLLAPSE and nodes 05/10 land almost on top of each other. Sounded together they beat, rather than forming an interval; §31.6 treats that as an acoustic design decision, not a defect.
3. **Node 10 FIXEDPOINT (index 10) and node 05 (index 5) are the same class.** `10 = 2 * 5`, so they are exactly two octaves apart. Since Node 10 FIXEDPOINT is Paper E Theorem 3.5 — the steady state the recursion converges to — and node 05 sits at the same class two octaves below, **the ordering places the fixed point in octave relation to an earlier node,**

which is precisely the relationship an accumulator has with its own transient. This is offered as an observation about the ordering's arithmetic, not as evidence for the ordering. Had the nodes been numbered differently the relation would not hold, and the sensitivity test of §29.9 would report exactly that.

Sensitivity test (ES-4), run. Three plausible re-orderings:

- Swap 09 and 10: class 9 <-> class 5, pitch-class displacement `772.63 − 407.82 = 364.81` cents. **Load-bearing.**
- Swap 03 and 12: classes both 3, displacement `0.00` cents. Harmonically invisible.
- Swap 07 and 11: class 7 <-> class 11, displacement `1102.64 − 737.65 = 364.99` cents. **Load-bearing.**

Maximum displacement 364.99 cents, well over the 100-cent threshold of ES-4. **Conclusion: the node ordering is load-bearing for this essence set, and any published tuning derived from it must carry the ordinal justification alongside it.** In this case the justification is available — the node order is the order of the method — and so the set is conformant. Had it not been available, ES-4 would have required publishing a range instead of a frequency.

Closing clause (ES-9). *Every value in both parts of this worked example is an essence frequency. It is a pitch assigned to a position in an ordered category set. It is not a physical resonance and implies none. No element, no node, no operator and no room physically resonates at any of these values, and this manual makes no claim that any does.*

Chapter 30 asks the question this chapter has deliberately left open: what happens when the two essences you want to relate do not share an ordering at all.

Chapter 30 . The Interdisciplinary Bridge

Chapter 29 left a hole in the middle of the method, and did so deliberately.

The essence law needs an index. Within a discipline the index is available: the categories are ordered, the ordering is defensible, and n is a position in that ordering. Between two disciplines the index is *not* available, and no amount of care in constructing either ordering supplies it. Biology's third category and Ethics' third category are both "third", and their threeness has nothing in common. Feed both into `f(n) = k n^2` and you get the same frequency for two things that have no relationship at all. That is not a small inconvenience. **It means the essence law, as published, cannot say anything about a relationship that crosses a category boundary — which is to say it cannot say anything about the only relationships this book has spent twenty-eight chapters learning to engineer.**

This chapter closes the hole. It closes it not with an approximation, a heuristic, or a convention, but with an identity that was already sitting in Paper A waiting to be used, and it closes it in one line:

```
n_eff = N_eff = 1 + (N − 1)u = e^F
```

The effective count IS the cross-disciplinary essence index. It is a real number rather than an integer, it is defined for any pair of states in the model, it equals 1 exactly when the two states are parallel — the ground essence — and it equals N exactly when they are orthogonal. Those are precisely the properties an index needs. Substituting it into the essence law produces the load bearing result of this entire volume, which is stated in §30.4 and which every subsequent section elaborates:

One nat of interaction surplus is exactly 2.885 octaves. The interaction surplus IS a pitch interval.

This is the keystone chapter. Chapters 1 through 28 built the arithmetic of interaction; Chapter 29 built the arithmetic of essence; this chapter shows they are the same arithmetic in two notations, and gives the exact conversion factor between them.

The Law

30.1 Statement of the bridge

IDENTIFICATION (badged). Let two essences sit in different blocks of Paper A's decomposition. Let `u = ||x ^ y||^2 = 1 − (x . y)^2` be the represented difference between their normalised state vectors, N the number of categories in the chosen partition, and `F = f(u) = ln(1 + (N−1)u)` the interaction surplus (Paper A, Theorem 2.1). Then the cross-disciplinary essence index is taken to be the effective count,

```
n_eff := N_eff = 1 + (N − 1)u = e^F
```

and the cross-disciplinary essence frequency follows from Chapter 29's law without adjustment:

```
f_cross = k ∗ n_eff^2 = k ∗ N_eff^2 = k ∗ e^(2F)

k = 0.483425732988414 Hz        (Chapter 29, §29.1)
```

Taking the base-2 logarithm gives the interval above the ground essence k:

```
OCTAVES ABOVE BASE  =  log2( f_cross / k )  =  log2( e^(2F) )  =  2F / ln 2

2 / ln 2 = 2.885390081777927        (computed)

OCTAVES  =  2.885390081777927 * F        with F in nats
```

Everything after the identification is exact algebra. The identification itself is a modeling decision and is badged as one. Its warrant is §30.3: `n_eff` is the unique quantity already present in the framework that satisfies every requirement an essence index must satisfy.

30.2 The honesty clauses, restated

The bridge inherits every clause from both parents. They are restated here in full because CANON requires it and because a chapter that joins two frameworks is exactly where an unstated clause would do the most damage.

From Chapter 29 (the essence law):

- **H1.** The output is an essence frequency, **not a physical resonance.** No claim is made that any material, molecule, element, institution, room or person physically resonates at these values, and the text must never imply otherwise.
- **H2.** k is a chosen reference constant, not a measured one. Its form — a Pythagorean interval over the golden section — is a design decision. **It cancels from every relationship in the system (Result One) and sets only the reference pitch.**
- **H3.** The square law is a postulate about essence structure. Its formal agreement with infinite-square-well eigenvalue spacing is a mathematical coincidence, not evidence of mechanism.
- **H4.** The essence index requires a defensible ordering. Within a discipline that means Paper B §1.2's validity conditions plus a declared ordinal rule. **Across disciplines the ordering problem does not go away — it becomes the requirement that u be honestly measured.**
- **H5 — the errata clause, on the strength of §29.4 and §29.10, binding on every chapter that quotes a number from the 2022 source.** Do not copy published values. Recompute them. `f = k n^2` reproduces all 118 published element values to a worst relative deviation of **9.0e-15**, and the compound rule `f = k^2 G^4` — `G` the count-weighted geometric mean of the constituent indices — matches **37 of the 44** published composites to better than one part in 10^12. **The seven that disagree are data-entry errors in the source, not a different law, and the O2 row proves it: for a single-element species `G` is the element index whatever the atom count, so no atom count whatever can produce the published figure.** Two transcription faults travel with the citation: the source's prose instruction `1+5^.5/2` is missing its parentheses and must be read as `(1 + sqrt(5))/2`, and the spreadsheet stored `phi` to fifteen significant figures, which accounts for the whole of the residual bias in the published table. **Cite the errata whenever the source is cited.** Where this chapter anchors to a published value — the `f(100)` check in §30.5 is the case — the anchor is to the recomputed value, and the agreement to fourteen figures is what licenses the check.

From Papers A through E (the surplus framework):

- **The S3 axiom is a design axiom.** `g(u) = e^{f(u)}` affine is *chosen*, not derived. The entire bridge rests on it, because `N_eff` is defined by it. Say so whenever you lean on it —

and this chapter leans on it in every section.

- **The triangle inequality FAILS for geometric triples** (Paper A §5). Subadditivity holds on the parameter u only. Do not chain intervals across three essences as though pitch distances added.
- **As deployed, the surplus functional is a LOWER BOUND** on the ensemble's information, not the information itself. We log the floor, not the ceiling. Therefore **every frequency this chapter computes is a lower bound on the pitch, not the pitch.**
- **Paper A is pure mathematics and makes no empirical or metaphysical claim.**
- **Paper D:** surplus is not identified with Shannon entropy or mutual information in general.
- **Paper E:** no thermodynamic negentropy law is proved. Higher static surplus does not imply higher canonical concentration.

One clause specific to the bridge, and it is the sharpest of them: the essence frequency computed here is a function of a *measurement of a relationship*, not of either party to it. `f_cross` changes when the measured u changes, and u changes with the encoder, the category scheme, the observer, and the day. **A cross-disciplinary essence frequency is a reading of an instrument, and like any instrument reading it carries the instrument's error, not the world's truth.**

The Engineering

30.3 Why there is no shared ordinal, and what would fix it

Paper A works in `V = R^(NM)`, decomposed into `N` orthogonal blocks `H_1, ..., H_N`, each of dimension `M`. A block is "an idealised category in a chosen representation" (Paper B §1.1); a unit vector `x` in `H_i` is "a normalised state-description within category `H_i`".

Within one block the essence law is comfortable. The block's `M` sub-modes have an ordering — Paper B's twelve-house example orders them Origin, Structure, Flow, Perception, Reciprocity, Transformation, Harmony, Growth, Memory, Repair, Connection, Completion — and if that ordering satisfies Chapter 29's §29.9 test then `n` is well defined and `f(n) = k n^2` applies.

Across blocks it is not. Consider two essences, `x` in `H_3` (Biology & Body) and `y` in `H_12` (Ethics & Governance). Each has an ordinal position *within its own block*. Neither position means anything relative to the other. Concretely, four distinct failures block the naive approach:

1. **No common origin.** Each block's index 1 is an arbitrary starting point of that block's ordering. There is no fact of the matter about whether Biology's first sub-mode "precedes" Ethics' first sub-mode.
2. **Index collision.** Two unrelated essences at position 3 in their blocks receive identical frequencies. **A method that assigns the same pitch to two things precisely because they are unrelated has inverted its own purpose.**
3. **No relational content.** The essence law as published takes one argument. It maps a category to a frequency. It has no slot for "the relationship between two categories", which is the object we actually care about.
4. **Scale mismatch.** Blocks may have different `M`. An index of 7 in a twelve-mode block and an index of 7 in a four-mode block are not comparable even in a loose sense.

What we need is a quantity that: (a) takes *two* states as arguments; (b) is a single positive real number, so it can be squared; (c) equals 1 for identical or representationally equivalent states, so that a relationship of no difference lands on the ground essence `k * 1^2 = k`; (d) is bounded above by a stated maximum, so the pitch is bounded; (e) is defined regardless of which blocks the two states occupy; and (f) is monotone in difference, so more difference means a higher index and therefore a higher pitch.

There is exactly one such quantity already in the framework, and it is not a coincidence that there is exactly one — the properties above are, item for item, the axioms S1 through S4 read backwards.

30.4 The effective count is the index

Paper A, Theorem 2.1: under axioms S1 (geometric dependence), S2 (`f(0) = 0`), S3 (affine effective count), and S4 (`f(1) = ln N`), the surplus functional is uniquely

```
F(x,y) = f(u) = ln(1 + (N−1)u),        u = 1 − (x . y)^2 in [0,1]
```

and the effective count is

```
N_eff = g(u) = e^{f(u)} = 1 + (N − 1)u
```

Check the six requirements against it, one by one:

requirement	`N_eff = 1 + (N−1)u`	source
(a) takes two states	yes, via `u(x,y)`	S1
(b) single positive real	yes, `N_eff >= 1 > 0`	affine in `u >= 0`
(c) equals 1 at no difference	`u = 0 –> N_eff = 1` exactly	S2, Thm 4.3
(d) bounded above	`u = 1 –> N_eff = N` exactly	S4, Thm 4.3
(e) block-independent	u is defined for any pair of unit vectors	Paper A §1
(f) monotone in difference	strictly increasing in u	Thm 3.1

Every box is ticked, and ticked by a theorem rather than by assertion. **The effective count is not merely *a* candidate for the cross-disciplinary essence index; it is the unique quantity the framework already contains that satisfies the requirements, and the requirements were derived from what an essence index must do, independently of the surplus framework.**

Its interpretation is also exactly right for the job. `N_eff` is a headcount *after the headcount has been audited for redundancy* (Chapter 7). Ten parallel windows are one window: `u = 0`, `N_eff = 1`, `F = ln 1 = 0`. That is the ground essence, `f = k * 1 = 0.483426 Hz`. **Two essences that are representationally the same thing produce the ground tone and no interval at all — which is the musically correct answer, because there is no interval between a note and itself.**

Substituting and simplifying:

```
f_cross  =  k * N_eff^2  =  k * (1 + (N−1)u)^2  =  k * e^(2F)
```

and, since `F = ln N_eff`,

```
log2(f_cross / k)  =  log2(N_eff^2)  =  2 * log2(N_eff)  =  2 * F / ln 2
```

Computed: `2 / ln 2 = 2.885390081777927`. One nat of interaction surplus is exactly 2.885390 octaves — equivalently 3462.4681 cents. And inversely, one octave of essence pitch is exactly `ln(2)/2 = 0.346574` nats of surplus.

Two verifications against values fixed independently in CANON_SONIC §4:

```
F = 0        (parallel, u = 0)
  -> N_eff = 1
  -> f_cross = k * 1 = 0.483425732988414 Hz     <- the ground essence
  -> octaves above base = 2 * 0 / ln 2 = 0.000000

F = ln 12 = 2.4849066497880004     (orthogonal at N = 12, u = 1)
  -> N_eff = 12.000000
  -> f_cross = k * 144 = 69.61330555033166 Hz
  -> octaves above base = 2 * 2.4849066497880004 / 0.6931471805599453
                        = 7.169925001442313
-> check: 2 * log2(12) = 2 * 3.584962500721156 = 7.169925001442313  [identical]
```

Both exact. **This is the unification: the ordinal index within a discipline and the effective count across disciplines are the same variable `n`. The two frameworks compose without adjustment, without a fitting constant, and without a correction term.**

It is worth pausing on how unusual that is. Nothing was tuned. `k` came from a 2022 document about chemistry; `N_eff` came from an axiomatic characterisation theorem about interaction geometry; the factor `2/ln 2` fell out of substituting one into the other. **The bridge has no free parameters at all** — which also means it has nothing to hide behind, and §30.11 states exactly how it fails.

30.5 The master conversion table

Everything an engineer needs from this chapter is in one table, computed here at full precision. Read it in either direction: from a measured surplus to a pitch, or from a chosen pitch to the surplus it implies.

F (nats)	`N_eff = e^F`	`f_cross = k e^(2F)` (Hz)	octaves above k	cents above k
0.000000	1.000000	0.483426	0.000000	0.0
0.050000	1.051271	0.534268	0.144270	173.1
0.100000	1.105171	0.590458	0.288539	346.2
0.200000	1.221403	0.721186	0.577078	692.5
0.300000	1.349859	0.880859	0.865617	1038.7
0.400000	1.491825	1.075884	1.154156	1385.0
0.500000	1.648721	1.314087	1.442695	1731.2
0.693147 (= ln 2)	2.000000	1.933703	2.000000	2400.0
0.750000	2.117000	2.166564	2.164043	2596.9
1.000000	2.718282	3.572060	**2.885390**	3462.5
1.250000	3.490343	5.889331	3.606738	4328.1
1.500000	4.481689	9.709865	4.328085	5193.7

F (nats)	N_eff = e^F	f_cross = k e^(2F) (Hz)	octaves above k	cents above k
1.750000	5.754603	16.008862	5.049433	6059.3
2.000000	7.389056	26.394151	5.770780	6924.9
2.197225 (= ln 9)	9.000000	39.157484	6.339850	7607.8
2.484907 (= ln 12)	12.000000	69.613306	7.169925	8603.9
2.750000	15.642632	118.290377	7.934823	9521.8
3.000000	20.085537	195.027860	8.656170	10387.4
3.218876 (= ln 25)	25.000000	302.141083	9.287712	11145.3
3.500000	33.115452	530.140688	10.098865	12118.6
4.000000	54.598150	1441.071800	11.541560	13849.9
4.605170 (= ln 100)	100.000000	4834.257330	13.287712	15945.3

Three checks that the table is internally consistent, each verifiable by eye:

- At `F = ln 2`, `N_eff = 2` and the pitch is exactly two octaves above base — Chapter 29's Lemma (doubling the index is two octaves) reappearing as a statement about surplus.
- At `F = ln 9` the frequency is `39.157484 Hz`, which is `k * 81` — identical to `f(9)` in Chapter 29's within-discipline table. **The bridge is continuous with the ordinal law: when `N_eff` lands on an integer, the cross-disciplinary frequency *is* the within-discipline frequency of that index.**
- At `F = ln 100`, `f = 4834.257330 Hz`, which is `f(100)` from the periodic-table verification of §29.4, matching the source document's published value for fermium to fourteen figures.

That second check deserves emphasis. It is the reason this is a bridge and not an analogy. **Whenever the audited effective count happens to equal an integer `m`, the cross-disciplinary essence frequency is exactly the within-discipline essence frequency of index `m`. The two formalisms agree on their overlap.** A bridge that disagreed with either shore would be a different construction with the same name.

The integer-anchored version of the same table, which is the one to use when you are working backwards from a target effective count:

N_eff	F = ln N_eff	f = k N_eff^2 (Hz)	octaves above k
1	0.000000	0.483426	0.000000
2	0.693147	1.933703	2.000000
3	1.098612	4.350832	3.169925
4	1.386294	7.734812	4.000000
5	1.609438	12.085643	4.643856
6	1.791759	17.403326	5.169925
7	1.945910	23.687861	5.614710
8	2.079442	30.939247	6.000000
9	2.197225	39.157484	6.339850
10	2.302585	48.342573	6.643856

N_eff	F = ln N_eff	f = k N_eff^2 (Hz)	octaves above k
11	2.397895	58.494514	6.918863
12	2.484907	69.613306	7.169925
16	2.772589	123.756988	8.000000
20	2.995732	193.370293	8.643856
24	3.178054	278.453222	9.169925
32	3.465736	495.027951	10.000000
50	3.912023	1208.564332	11.287712
100	4.605170	4834.257330	13.287712

30.6 Reading the bridge backwards: pitch as a unit of surplus

The conversion runs both ways, and the reverse direction is the one that changes how you think about the numbers. Since `F = (cents/1200) * ln(2) / 2 = cents * ln(2) / 2400`:

interval	cents	F (nats)	N_eff = e^F	u required at N = 12
unison	0	0.000000	1.000000	0.000000
semitone	100	0.028881	1.029302	0.002664
whole tone	200	0.057762	1.059463	0.005406
minor 3rd	300	0.086643	1.090508	0.008228
major 3rd	400	0.115525	1.122462	0.011133
perfect 4th	500	0.144406	1.155353	0.014123
tritone	600	0.173287	1.189207	0.017201
perfect 5th	700	0.202168	1.224054	0.020369
minor 6th	800	0.231049	1.259921	0.023629
major 6th	900	0.259930	1.296840	0.026985
minor 7th	1000	0.288811	1.334840	0.030440
major 7th	1100	0.317692	1.373954	0.033996
octave	1200	0.346574	1.414214	0.037656
two octaves	2400	0.693147	2.000000	0.090909
three octaves	3600	1.039721	2.828427	0.166221

Two facts in that table are worth carrying in the head.

A perfect fifth of surplus is 0.202168 nats. That is the whole content of the sentence "the interaction surplus IS a pitch interval": intervals are now a unit of measurement for organisational difference, and a familiar one. When Chapter 9 says a rotation programme raised u from 0.28 to 0.41, that is now sayable as *the group's cross-disciplinary tone rose by 1040 cents* — a shade under an octave. It is the same statement. It is easier to feel.

And the second fact, which is a warning: the entire audible range is about 3.4 nats. Ten octaves — 20 Hz to 20 kHz, roughly the human hearing range — is `10 / 2.885390 = 3.466` nats. `ln N` for `N = 12` is 2.485 nats, so **a twelve-category scheme's entire surplus range, from parallel to orthogonal, spans 7.17 octaves: about two-thirds of everything a human can hear.** For `N = 50` the span is `2 * log2(50) = 11.29` octaves — more than the audible

range. Beyond about `N = 32` the essence pitch of an orthogonal pair leaves the audible band entirely and must be folded. §30.9 and Chapter 31 handle the folding.

The corresponding ceiling table, since `F <= ln N` is hard (Paper D, Corollary 2.2):

N	`ln N` (max surplus)	max octaves = 2 `log2 N`	max `f = k N^2` (Hz)
2	0.693147	2.000000	1.9337
3	1.098612	3.169925	4.3508
5	1.609438	4.643856	12.0856
7	1.945910	5.614710	23.6879
10	2.302585	6.643856	48.3426
12	2.484907	7.169925	69.6133
20	2.995732	8.643856	193.3703
50	3.912023	11.287712	1208.5643
100	4.605170	13.287712	4834.2573
118	4.770685	13.765286	6731.2199

`ln N` is the hard ceiling on surplus per interaction (Paper D, Cor 2.2), and it is therefore a hard ceiling on pitch: `2 log2 N` octaves above the ground essence, and not one cent more. The table's last row is the periodic table's own ceiling and lands on 6731.2199 Hz — the source document's published value for oganesson, index 118, to fourteen figures. That agreement is not a new fact; it is §30.5's continuity property showing up at the top of the range as well as the middle.

Chapter 8's lesson is audible here. Doubling `N` adds `ln 2 = 0.693147` nats, hence exactly two octaves, no matter where you start. **You cannot escape a headroom problem by adding domains, and now you cannot escape a pitch problem that way either: going from twelve categories to twenty-four buys you two octaves and nothing else.**

30.7 Crossing and divergence: Theorem 4.1 in pitch

Paper A's Theorem 4.1 splits u into two sources, and the split survives translation into pitch with its bounds intact.

For a pure block vector `x` in `H_i` and any unit `y`, write `y = y_par + y_perp` with `y_par` in `H_i` and `y_perp` orthogonal to `H_i`. Let `alpha = ||y_par||`, `beta = ||y_perp||`, and `gamma = angle(x, y_par)`. Then

```
    u = beta^2 + alpha^2 sin^2(gamma)  =  u_cross + u_div

     u_cross = beta^2                        BETWEEN categories   (cross-
disciplinary difference)
    u_div   = alpha^2 sin^2(gamma)    WITHIN a category    (internal dissent)

    max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

Multiply the bound through by `2/ln 2 = 2.885390` — a positive constant, so the inequalities are preserved exactly:

```
max( OCT(u_cross), OCT(u_div) )  <=  OCT(u)  <=  OCT(u_cross) + OCT(u_div)

 where OCT(v) := 2.885390081777927 * ln(1 + (N-1)v)
```

In musical terms: the cross-disciplinary tone of a relationship is never lower than the higher of its two component tones, and never higher than the two component tones stacked. The lower bound says you always get at least the better of your two sources of difference. The upper bound says crossing and divergence never combine to give you more than they would separately — **there is no synergy beyond additivity, and the bound is the framework refusing to promise any.**

Computed, for `N = 12`, showing the whole structure:

u_cross	u_div	u	f(u_cross)	f(u_div)	f(u)	max bound	sum bound	f_cross (Hz)	octaves
0.00	0.00	0.00	0.000000	0.000000	0.000000	0.000000	0.000000	0.483426	0.0000
0.20	0.00	0.20	1.163151	0.000000	1.163151	1.163151	1.163151	4.950280	3.3561
0.00	0.20	0.20	0.000000	1.163151	1.163151	1.163151	1.163151	4.950280	3.3561
0.10	0.10	0.20	0.741937	0.741937	1.163151	0.741937	1.483875	4.950280	3.3561
0.40	0.10	0.50	1.686399	0.741937	1.871802	1.686399	2.428336	20.424737	5.4009
0.10	0.40	0.50	0.741937	1.686399	1.871802	1.686399	2.428336	20.424737	5.4009
0.25	0.25	0.50	1.321756	1.321756	1.871802	1.321756	2.643512	20.424737	5.4009
0.45	0.45	0.90	1.783391	1.783391	2.388763	1.783391	3.566782	57.435816	6.8925
0.90	0.00	0.90	2.388763	0.000000	2.388763	2.388763	2.388763	57.435816	6.8925
0.50	0.50	1.00	1.871802	1.871802	2.484907	1.871802	3.743604	69.613306	7.1699
0.80	0.20	1.00	2.282382	1.163151	2.484907	2.282382	3.445533	69.613306	7.1699

And the same in octaves, with the slack — the gap between the achieved pitch and the sum bound — made explicit:

u_cross	u_div	achieved OCT(u)	max bound (oct)	sum bound (oct)	slack (oct)
0.20	0.00	3.3561	3.3561	3.3561	0.0000
0.00	0.20	3.3561	3.3561	3.3561	0.0000
0.10	0.10	3.3561	2.1408	4.2816	0.9254
0.40	0.10	5.4009	4.8659	7.0067	1.6058
0.25	0.25	5.4009	3.8138	7.6276	2.2267
0.45	0.45	6.8925	5.1458	10.2916	3.3990
0.50	0.50	7.1699	5.4009	10.8018	3.6318
0.90	0.00	6.8925	6.8925	6.8925	0.0000

Read the slack column and the structure of the bound becomes visible. **Slack is zero exactly when one of the two sources is zero, and it is maximal when the two are balanced.** At `u_cross = u_div = 0.5` the sum bound promises 10.80 octaves and the system delivers 7.17 — a slack of 3.63 octaves, more than a third of the audible range, thrown away by concavity.

That is not a defect. **It is strict concavity (Paper A, Theorem 3.3) charging its price, and the price is real: two half-sized sources of difference produce a lower tone than one full-sized one.** Compare rows: `u_cross = 0.2, u_div = 0` and `u_cross = 0.1, u_div = 0.1` both give `u = 0.2` and therefore both give 3.3561 octaves — identical pitch, because `f` depends only on the sum. But the *bound* differs: the split case had 4.28 octaves available under the sum bound and reached 3.36. The concentrated case had 3.36 available and reached 3.36, saturating its bound.

The design lesson is Chapter 10's, now audible: **crossing and divergence are not substitutes, and a system rich in one and starved of the other sits at exactly the same pitch as one with the same total u distributed differently — but it has a very different fragility.** If your entire u is `u_cross` (cross-domain hires, one permitted opinion inside each domain) then a single reorganisation that merges two domains destroys a large fraction of your surplus in one move. If your u is split, no single structural change takes it all. **The pitch does not distinguish these cases; the decomposition does. Report both.**

30.8 What a two-essence interval sounds like as u sweeps 0 to 1

The most useful intuition in this chapter comes from watching one number move. Fix `N = 12`. Sweep u from 0 (representationally identical) to 1 (orthogonal) and watch what happens to the tone relative to the ground essence k.

u	`N_eff = 1 + 11u`	F (nats)	`f_cross` (Hz)	octaves above k	folded into [220,440) Hz
0.00	1.0000	0.000000	0.483426	0.000000	247.5140 (j = +9)
0.05	1.5500	0.438255	1.161430	1.264536	297.3262 (j = +8)
0.0909 (= 1/11)	2.0000	0.693147	1.933703	2.000000	247.5140 (j = +7)
0.10	2.1000	0.741937	2.131907	2.140779	272.8842 (j = +7)
0.15	2.6500	0.974560	3.394857	2.811985	434.5417 (j = +7)
0.20	3.2000	1.163151	4.950280	3.356144	316.8179 (j = +6)
0.25	3.7500	1.321756	6.798174	3.813781	435.0832 (j = +6)
0.30	4.3000	1.458615	8.938542	4.208673	286.0333 (j = +5)
0.40	5.4000	1.686399	14.096694	4.865919	225.5471 (j = +4)
0.50	6.5000	1.871802	20.424737	5.400879	326.7958 (j = +4)
0.60	7.6000	2.028148	27.922670	5.851999	223.3814 (j = +3)
0.70	8.7000	2.163323	36.590494	6.242031	292.7239 (j = +3)

u	N_eff = 1 + 11u	F (nats)	f_cross (Hz)	octaves above k	folded into [220,440) Hz
0.75	9.2500	2.224624	41.363114	6.418907	330.9049 (j = +3)
0.80	9.8000	2.282382	46.428207	6.585563	371.4257 (j = +3)
0.90	10.9000	2.388763	57.435811	6.892512	229.7432 (j = +2)
1.00	12.0000	2.484907	69.613306	7.169925	278.4532 (j = +2)

And the interval itself — the ratio N_eff^2 between the pair's tone and the ground essence, expressed as a musician would hear it:

u	ratio N_eff^2	cents	folded cents	nearest 12-TET
0.00	1.000000	0.00	0.00	unison +0.0
0.02	1.488400	688.51	688.51	perfect 5th -11.5
0.05	2.402500	1517.44	317.44	minor 3rd +17.4
0.10	4.410000	2568.93	168.93	major 2nd -31.1
0.15	7.022500	3374.38	974.38	minor 7th -25.6
0.20	10.240000	4027.37	427.37	major 3rd +27.4
0.25	14.062500	4576.54	976.54	minor 7th -23.5
0.30	18.490000	5050.41	250.41	minor 3rd -49.6
0.40	29.160000	5839.10	1039.10	minor 7th +39.1
0.50	42.250000	6481.06	481.06	perfect 4th -18.9
0.60	57.760000	7022.40	1022.40	minor 7th +22.4
0.70	75.690000	7490.44	290.44	minor 3rd -9.6
0.80	96.040000	7902.68	702.68	perfect 5th +2.7
0.90	118.810000	8271.01	1071.01	major 7th -29.0
1.00	144.000000	8603.91	203.91	major 2nd +3.9

Four things to take from the sweep, each an engineering fact rather than an aesthetic one.

(1) The rise is violently front-loaded. Concavity, made audible. Octaves bought per 0.1 of u:

u range	octaves gained
0.0 -> 0.1	2.1408
0.1 -> 0.2	1.2154
0.2 -> 0.3	0.8525
0.3 -> 0.4	0.6572
0.4 -> 0.5	0.5350
0.5 -> 0.6	0.4511
0.6 -> 0.7	0.3900
0.7 -> 0.8	0.3435

u range	octaves gained
0.8 -> 0.9	0.3069
0.9 -> 1.0	0.2774

The first tenth of u buys 2.14 octaves; the last tenth buys 0.28 — a factor of 7.7. **Chapter 7's merciful clause, heard: the second genuinely different member is the largest single gain you will ever buy, and in pitch terms it is more than two octaves' worth.** A team at u = 0.05 that reaches u = 0.15 gains 1.55 octaves. The same team going from u = 0.85 to u = 0.95 gains 0.29. The cheap improvements are at the bottom.

(2) The sweep is not monotone once folded. In the raw column, pitch rises monotonically with u — guaranteed by Theorem 3.1. In the folded column it does not: u = 0.15 folds to 434.54 Hz and u = 0.20 folds to 316.82 Hz, which is *lower*. Folding wraps. **This is the single most common error in reading these tables and it must be stated every time: a folded frequency is a pitch class, not a pitch, and pitch classes do not order.** If you need to compare two systems by surplus, compare F, or compare raw f_cross, or compare octaves-above-base. Never compare folded values.

(3) At u = 1/11 = 0.0909... the effective count is exactly 2 and the interval is exactly two octaves. This is the smallest u at N = 12 that "buys a second full member", and it is worth knowing as a landmark: below u = 0.091, twelve categories are worth fewer than two.

(4) The extremes are the two facts you will quote most. At u = 0 the two essences are the same essence and the interval is a unison — the pair sounds as one note at the ground essence 0.483426 Hz. At u = 1 the pair is orthogonal, N_eff = 12, the interval is 144:1 — seven octaves and a Pythagorean whole tone, 8603.91 cents. **Perfect orthogonality at twelve categories is a major second, seven octaves up. That is the whole dynamic range of the instrument.**

30.9 Sensitivity: how precisely must u be measured?

Because the bridge has no free parameters, all of its uncertainty lives in u. Paper A, Theorem 4.2 gives the sharp Lipschitz constant N − 1, so

```
dF/du = (N−1) / (1 + (N−1)u)        maximum N−1, attained at u = 0

d(octaves)/du = 2.885390 * dF/du
```

N	max dF/du (at u=0)	max d(oct)/du	dF/du at u=1	d(oct)/du at u=1
2	1.0000	2.8854	0.5000	1.4427
5	4.0000	11.5416	0.8000	2.3083
12	11.0000	31.7393	0.9167	2.6449
20	19.0000	54.8224	0.9500	2.7411
50	49.0000	141.3841	0.9800	2.8277

At N = 12 and u near zero the pitch moves 31.74 octaves per unit of u. In practical terms, with a just-noticeable pitch difference of about 5 cents:

u	dF/du	d(cents)/du	du producing a 5-cent shift
0.00	11.0000	38087.1	0.000131
0.10	5.2381	18136.7	0.000276
0.25	2.9333	10156.6	0.000492
0.50	1.6923	5859.6	0.000853
0.75	1.1892	4117.5	0.001214
1.00	0.9167	3173.9	0.001575

A change in u of 0.00013 — one part in seven thousand — is audible as a pitch shift near u = 0. No organisational measurement of u is anywhere near that precise. Chapter 24's estimation methods (disagreement rate, error decorrelation, prediction diversity) yield u to perhaps two decimal places on a good day, and the second of those is soft.

This has one hard consequence and it is a specification, not a caution:

Publish u with an interval, and publish the pitch as the interval's image. If your survey supports u = 0.28 +/− 0.05, then at N = 12 the tone is not 8.0473 Hz; it is the range from k∗(1 + 11∗0.23)^2 = 6.0239 Hz to k∗(1 + 11∗0.33)^2 = 10.3631 Hz, which is 0.7827 octaves wide — a minor sixth of uncertainty. **Quote the sixth. A single-valued essence frequency from a two-significant-figure u is precision theatre (FM-29.6) wearing a new hat.**

The second half of the sensitivity question is temporal, and it has an exact answer. The table above asks how small a difference in u is audible. It does not ask how *long* the tone must sound before the difference can be heard at all, and that question is not a matter of taste: the Gabor limit fixes it. With the time-bandwidth product dt ∗ df >= 1/(4 pi) and the bridge's own relation d(ln f) = 2 dF, Chapter 71 derives the bound

```
cycles required  >=  1 / (8 pi dF)
```

Resolving a surplus difference dF requires at least 1/(8 pi dF) cycles of the rendered tone, and the bound is independent of the transposition — because it is a bound on a ratio, and octave transposition multiplies numerator and denominator alike. **This is the exact Gabor bridge to Paper A: it prices, in time, what it costs to hear a difference in surplus.**

dF (nats)	cycles required	duration at 69.613 Hz (N_eff = 12)	at 302.141 Hz (N_eff = 25)
0.20	0.199	2.86 ms	0.66 ms
0.10	0.398	5.72 ms	1.32 ms
0.05	0.796	11.43 ms	2.63 ms
0.02	1.989	28.58 ms	6.58 ms
0.01	3.979	57.16 ms	13.17 ms
0.005	7.958	114.31 ms	26.34 ms
0.001	39.789	571.57 ms	131.69 ms

Read the two sensitivity results together and they close on each other from opposite sides. The Lipschitz table says a du of 0.00013 is audible in principle near u = 0; the Gabor bound says that hearing it requires the tone to persist. **Neither one alone is a specification. Together they**

say: the pitch difference your measurement can honestly support is coarse, and coarse differences are cheap in time — which is the fortunate direction for the error to run. A rendering that sounds each essence for well under a cycle of the resolution it claims is asserting a distinction that physics has not delivered, whatever the frequency table says. Chapter 71 develops the full set of design laws that follow from the limit; Chapter 31 applies them to built work.

30.10 Composing with the storage law

The bridge produces a pitch from `u`. Chapter 6's headroom gauge produces a budget from `Q`, `delta`, `eta`, `C` and `N`. They meet in `ln N`, and it is worth being explicit about how, because a reader will otherwise try to make the pitch do work it cannot do.

```
h_t   =  eta * ln N - delta * Q_t - C_t            (headroom, Paper E §6.4)
 OCT   =  2.885390 * F,       F = ln N_eff <= ln N     (this chapter)
```

Both contain `ln`-of-a-count, and both are ceilinged by `ln N`. But they are different quantities and must never be added:

- `F` is the surplus of **one interaction**, and `OCT` is its pitch. It is an instantaneous, geometric quantity.
- `h_t` is a **budget over a period**, in units of stored order. It depends on the history `Q_t`.

The correct composition is the one Paper E already specifies: `S_t = F(U_t)` is the surplus *input* to the storage recursion. So the pitch is a gauge on the **input**, not on the **stock**:

```
OCT_t = 2.885390 * S_t      is the pitch of this period's interaction

Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t
```

The tone tells you what is going into the accumulator this period. It tells you nothing about how much is already in it. A group can be sounding a high tone — rich, genuinely non-parallel interaction — while `Q` collapses, because `eta` is near zero and none of it is sticking (Chapter 4's Leaky Vessel), or because `C` exceeds `eta ln N - delta Q` and Corollary 3.4 has already sealed the outcome. Chapter 31 uses the pitch as a design target for rooms and cadences; it does not use it as a health metric, and neither should you.

One legitimate joint reading, and it is useful: since `S_t <= ln N` and `OCT_t = 2.885390 S_t`, the maximum pitch a system can sound is fixed by its category count, and the maximum sustainable *increment* to stored order is `eta ln N - delta Q_t - C_t`. **Both ceilings come from the same `ln N`, which means a system that has run out of pitch has also run out of headroom growth from geometry — and the remaining levers are `eta`, `C`, and shedding `Q` (Chapter 6), none of which change the tone at all.** That is a genuinely useful diagnostic: *if the tone is already at ceiling and headroom is still negative, no amount of further diversity work will help, and continuing to do diversity work is the mistake.*

30.11 What the bridge does not claim

Stated flatly, because a keystone that is over-claimed takes the arch down with it.

1. **It does not claim that essences resonate.** H1. `f_cross` is an essence frequency of a *relationship*, computed from a *measurement* of that relationship. Nothing physical resonates.

2. **It does not claim that `2/ln 2` is a natural constant.** It is `2/ln 2` because Chapter 29's law is square and octaves are base two. Had the essence law been cubic, the factor would be `3/ln 2 = 4.328`. **The factor is a consequence of two design choices, and its exactness is the exactness of arithmetic, not of nature.**
3. **It does not claim to have derived `N_eff` as the index.** The identification is badged. What is *shown* is that `N_eff` uniquely satisfies six requirements (§30.4), that the substitution is exact, and that the result is continuous with the within-discipline law at every integer (§30.5). That is a strong warrant. It is not a derivation.
4. **It does not claim more information than `F` contains.** As deployed the surplus functional is a **lower bound** on the ensemble's information. Therefore `f_cross` is a **lower bound on the pitch**. If you must err, err by reporting a tone that is too low.
5. **It does not permit chaining.** The triangle inequality fails for geometric triples (Paper A §5). Given essences A, B, C, you may compute the A-B tone and the B-C tone, but **you may not add the two intervals to obtain the A-C tone.** Compute A-C directly from `u(A,C)` or do not compute it. This is the failure mode that will catch the most people, because pitch intervals *do* add in ordinary music and the habit is deeply worn.
6. **It does not survive an undisclosed change of `N`.** `N_eff = 1 + (N−1)u` contains `N` explicitly. Recount your categories and every published frequency moves. Paper B is explicit that `N = 12` "is not asserted as correct" and that other decompositions produce different maxima though the same structural theorems. **`N` must be published with every frequency, as `k` and `j` already are.**

The Specification

BR-1 — Publish `(N, u, F, N_eff, f_cross, j, band)` as an inseparable tuple. A cross-disciplinary essence frequency quoted without its `N` and its `u` is unreproducible and therefore not a result.

BR-2 — Badge the identification. State, at each first use in a document: *the effective count is taken as the cross-disciplinary essence index; this is a modeling identification, warranted by the six-requirement argument, not derived.*

BR-3 — Badge S3. The affine effective-count axiom is a design axiom, chosen not derived. `N_eff` exists because S3 was adopted. The bridge rests on S3 entirely; say so.

BR-4 — Publish `u` as an interval and the pitch as its image. Never publish a single-valued `f_cross` from a `u` estimated to two decimal places. Report the pitch range in octaves or cents. (§30.9)

BR-5 — Never compare folded frequencies to infer surplus ordering. Folding wraps and destroys order. Compare `F`, raw `f_cross`, or octaves-above-base. (§30.8, point 2)

BR-6 — Never add two intervals to obtain a third. The triangle inequality fails for geometric triples. Compute every pair directly from its own `u`. (§30.11, point 5)

BR-7 — Report the decomposition, not only the total. Publish `u_cross` and `u_div` separately with the bound `max(OCT(u_cross), OCT(u_div)) <= OCT(u) <= OCT(u_cross) +`

OCT(u_div) evaluated. Two systems at the same pitch with different decompositions have different fragilities. (§30.7)

BR-8 — State the ceiling with the value. Every published cross-disciplinary tone carries 2 log2 N octaves as its stated maximum. A reader must be able to see how much of the range is in use.

BR-9 — Do not read the tone as a health metric. OCT_t = 2.885390 S_t gauges the *input* to the storage recursion, not the stock Q_t, and not h_t. (§30.10)

BR-10 — Restate H1 and the lower-bound clause on any page carrying a frequency. "Essence frequency, not physical resonance" and "the surplus functional as deployed is a lower bound on the ensemble's information, so this tone is a lower bound on the pitch."

BR-11 — Re-derive when N changes. Any change to the category partition invalidates every published frequency in the set. Republish or withdraw; do not let a stale tone circulate.

BR-12 — Verify the integer-anchor property on any new implementation. For a test value of u making N_eff an integer m, confirm your code returns exactly k m^2 — the same number Chapter 29's within-discipline table gives for index m. This is a two-line regression test that catches almost every implementation error.

The Failure Modes

FM-30.1 THE FALSE ORDINAL ACROSS BLOCKS

Damages: every cross-disciplinary frequency computed. *Description.* Two essences in different blocks are given ordinals within their own blocks, and those ordinals are fed into f(n) = k n^2 as though comparable. Biology's third sub-mode and Ethics' third sub-mode receive the same frequency. *Signature.* Two unrelated essences with identical published tones, and the practitioner reporting this as a discovered affinity. *Fix.* The whole of this chapter. Use n_eff = 1 + (N–1)u. Within-block ordinals may be used within a block and nowhere else.

FM-30.2 THE CHAINED INTERVAL

Damages: any three-party analysis. *Description.* The A-B tone is a perfect fifth up, the B-C tone is a fourth up, therefore A-C is an octave up. **It is not.** The triangle inequality fails for geometric triples (Paper A §5). *Signature.* An interval reported for a pair whose u was never measured. *Fix.* BR-6. Measure u(A,C) or report nothing for that pair.

FM-30.3 THE UNAUDITED ROSTER

Damages: u, and therefore everything. *Description.* N_eff is computed from a headcount rather than from a measured u. Twelve people in a room are reported as N_eff = 12, tone at 69.61 Hz, when the measured u is 0.15 and the true N_eff is 2.65. *Signature.* N_eff equal or nearly equal to the literal headcount, with no disagreement audit in the file. *Fix.* Chapter 7 and Chapter 24. N_eff is a headcount **after audit for redundancy**. Measure u. An unmeasured u is not an optimistic estimate; it is no estimate.

FM-30.4 THE FOLDED COMPARISON

Damages: every ordering claim between systems. *Description.* System A folds to 434.54 Hz, system B to 316.82 Hz, therefore A has more surplus. It does not: those are `u = 0.15` and `u = 0.20` respectively, and B has more. *Signature.* A comparison between two folded values. *Fix.* BR-5. Compare in `F` or in octaves-above-base, always.

FM-30.5 PRECISION FROM A SOFT `u`

Damages: credibility, exactly as FM-29.6 does. *Description.* `u = 0.28` from a small disagreement audit becomes `f_cross = 8.0473 Hz`, published to four decimals. The underlying interval is 0.638 octaves wide. *Signature.* More precision in the output than in the input. *Fix.* BR-4. Publish the range. §30.9 gives the arithmetic.

FM-30.6 THE SILENT `N` CHANGE

Damages: comparability across time. *Description.* An organisation reorganises from twelve departments to nine. Nobody re-derives. Frequencies computed under `N = 12` are compared with new observations under `N = 9`, and the apparent change is entirely an artefact of the denominator. *Signature.* A step change in the reported tone that coincides with a reorganisation date. *Fix.* BR-11. `N` travels with every value. On any change to the partition, republish or withdraw.

FM-30.7 THE SYNERGY CLAIM

Damages: the subadditive bound, and any budget built on it. *Description.* "Combining cross-domain hiring with an internal red team will more than double our surplus." The bound says it will not: `f(u) <= f(u_cross) + f(u_div)`, always. *Signature.* A projected gain exceeding the sum of the parts' separate gains. *Fix.* BR-7. Evaluate the sum bound and cap the projection at it. Note that in practice you will land well *below* the sum bound — §30.7's slack column shows 3.6 octaves of concavity tax at balanced sources.

FM-30.8 THE INVERTED READING

Damages: the entire method, by turning it into decoration. *Description.* A frequency is chosen first — because it sounds pleasant, or because it is 432 Hz, or because a client asked for it — and u is then back-solved to justify it. *Signature.* A u value with more than two significant figures that was never measured, or a u that lands suspiciously close to a round musical interval. *Fix.* Chapter 25's preregistration discipline, applied without exception: **fix the measurement protocol and the ordinal rule before computing any frequency, and freeze them.** The tone is an output. If you want a particular tone, say so openly and choose it as a design decision — that is legitimate, and Chapter 31 does exactly that when it selects a band — but do not dress a chosen tone as a measured one.

FM-30.9 THE TONE AS SCOREBOARD

Damages: `Q`, via misdirected effort. *Description.* The cross-disciplinary tone is adopted as an organisational KPI. It rises. `Q` falls, because `eta` is near zero and `C` is above the sustainability ceiling. *Signature.* A rising tone alongside falling stored order; or, worse, a rising tone and no

measurement of stored order at all. *Fix.* BR-9, and Goodhart. **The tone gauges the input to the accumulator. Measure `h_t` separately and let it, not the tone, decide whether you may build.**

FM-30.10 THE UNBADGED IDENTIFICATION

Damages: the reader's ability to audit the argument. *Description.* `n_eff = N_eff` is presented as a theorem. It is not; it is a modeling identification with a strong six-part warrant. *Signature.* The sentence "it follows that the effective count is the essence index", with no badge. *Fix.* BR-2 and BR-3. And be precise about what *is* proved: the substitution's consequences, the factor `2/ln 2`, the integer-anchor continuity, and the transported Theorem 4.1 bounds are all exact. The identification is the postulate; everything downstream of it is arithmetic.

Worked Example

The Institute: a twelve-house research organisation, carried end to end

An interdisciplinary research institute is organised into the twelve houses of Paper B §2.1 — a partition drawn from the Vovina Ontological OmniTautology corpus, with `N = 12` and `M = 12` sub-modes per house. Paper B is explicit that twelve "is not asserted as correct" and that other decompositions produce different maxima with the same structural theorems; that caveat travels with every number below.

Step 0 — Validity check (BR-1, ES-5). Paper B §1.2 conditions:

- *State encoding.* Each researcher's working position is encoded as a normalised vector over the 144-dimensional space (12 houses x 12 sub-modes), from a structured self-description plus a publication-topic coding. Declared.
- *Category assignment.* The twelve houses. Declared. `N = 12`.
- *Geometric dependence.* Interaction is modelled through `u = 1 − (x.y)^2`, an inner-product-derived quantity. Satisfied.
- *Zero baseline.* Two researchers with representationally equivalent positions give `u = 0`. Satisfied.
- *Affine effective count.* Adopted as a modeling assumption. **Badged (BR-3).**
- *Normalisation.* Maximum surplus `ln 12 = 2.484907` nats. Declared.

All six stated. The domain is in scope.

Step 1 — Measure u. Chapter 24's protocol, run over four months: independent-estimate- before-discussion on 34 forecast questions, error decorrelation computed pairwise, disagreement rate logged from the decision log.

```
measured u (institute-wide mean over pairs) = 0.28,  supported to +/- 0.05
```

Step 2 — Compute the baseline tone.

```
N_eff = 1 + 11 * 0.28 = 4.0800
F     = ln 4.0800     = 1.406097 nats
OCT   = 2.885390081777927 * 1.406097 = 4.0571 octaves above k
f_cross = k * 4.08^2  = 0.483425732988414 * 16.6464 = 8.0473 Hz
```

```
folded into [220, 440): 8.0473 * 2^5 = 257.5135 Hz   (j = +5)
                        nearest note C4 (261.626 Hz), −27 cents
```

Step 3 — Publish the uncertainty as an interval (BR-4). With `u = 0.28 +/− 0.05`:

```
u = 0.23 -> N_eff = 3.5300 -> f_cross =  6.0239 Hz -> 3.6393 octaves above k
u = 0.33 -> N_eff = 4.6300 -> f_cross = 10.3631 Hz -> 4.4220 octaves above k

published tone: 8.05 Hz, range 6.02 – 10.36 Hz, i.e. 4.06 octaves above base
                with an uncertainty of 0.7827 octaves (a minor sixth)
```

That 0.78-octave band is the honest result. The four-decimal figure in Step 2 is an intermediate, not a finding.

Step 4 — Decompose (BR-7). The disagreement audit separates between-house disagreement from within-house disagreement:

```
u_cross = 0.20   (between houses)
u_div   = 0.08   (within houses)
u       = 0.28   [checks: 0.20 + 0.08 = 0.28, Paper A Thm 4.1]

f(u_cross) = ln(1 + 11*0.20) = 1.163151 nats -> 3.3561 octaves
f(u_div)   = ln(1 + 11*0.08) = 0.631272 nats -> 1.8215 octaves
f(u)       = ln(1 + 11*0.28) = 1.406097 nats -> 4.0571 octaves

bound: max(3.3561, 1.8215) = 3.3561 <= 4.0571 <= 3.3561 + 1.8215 = 5.1776
 slack to the sum bound: 5.1776 − 4.0571 = 1.1205 octaves
```

Diagnosis: the institute is 71% crossing and 29% divergence. It has bought cross-domain difference (that is what an interdisciplinary institute is for) and is thin on internal dissent — Chapter 10's classic pattern. The fragility follows: a merger of two houses would remove a substantial fraction of `u_cross`, and there is little `u_div` underneath to catch the fall.

Step 5 — Design the intervention (Chapter 26). Two moves, sequenced:

- *Rotation across houses* (raises `u_cross`): six-month placements, four researchers per cycle.
- *Standing red team* (raises `u_div`): rotating adversarial role inside each house, with a minority report requirement on every major decision.

Projected post-intervention `u = 0.41`, from the pilot's measured effect on eight pairs.

Step 6 — Compute the interval the intervention buys.

```
before: u = 0.28, N_eff = 4.0800, F = 1.406097
after:  u = 0.41, N_eff = 5.5100, F = 1.706565

dF = 1.706565 − 1.406097 = 0.300468 nats

interval ratio = (5.5100 / 4.0800)^2 = 1.823824
cents          = 1200 * log2(1.823824) = 1040.36
octaves        = 0.8670

cross-check via the bridge constant:
```

```
2.885390081777927 * 0.300468 = 0.866967 octaves      [identical to 0.8670]
```

The intervention is worth 1040 cents — a minor seventh, a shade under an octave. That is the entire yield of a two-year programme of rotation and red-teaming, stated in a unit anyone in the room can hear.

Post-intervention tone:

```
 f_cross = k * 5.51^2 = 14.6769 Hz
folded into [220,440): 14.6769 * 2^4 = 234.8297 Hz  (j = +4), A#3 +13 cents
 octaves above k: 4.9241
 fraction of the N = 12 ceiling used: 4.9241 / 7.169925 = 68.7%
```

Step 7 — State the ceiling (BR-8). `N = 12`, so the maximum is `2 log2 12 = 7.169925` octaves, reached only at `u = 1`. The institute at `u = 0.41` uses 68.7% of its available range. A third intervention taking u to 0.52 would yield:

```
u = 0.52 -> N_eff = 6.7200 -> F = 1.905088 -> 5.4969 octaves (76.7% of ceiling)
 incremental gain from 0.41 -> 0.52: 0.5728 octaves = 687 cents
```

Note the decay: u 0.28 -> 0.41 (a step of 0.13) bought 1040 cents; u 0.41 -> 0.52 (a step of 0.11) buys 687 cents. Concavity charging its price, exactly as §30.8's table predicts.

Step 8 — Do not read the tone as health (BR-9). The institute's storage parameters, measured separately per Chapter 24:

```
eta = 0.35     conversion of surplus into durable structure
delta = 0.08   quarterly decay of stored order
Q = 18.0       stored order, in the institute's declared units
C = 0.40       dissipation per quarter

h = eta * ln N - delta * Q - C
  = 0.35 * 2.484907 - 0.08 * 18.0 - 0.40
  = 0.869717 - 1.440000 - 0.400000
  = -0.9703
```

Headroom is negative. `h < 0` means unsustainable in this model: no choice of interaction geometry prevents order from decreasing (Paper E, Corollary 3.4). The tone is rising and the institute is losing stored order, and the two facts are entirely compatible. The maintenance term alone, `delta Q = 1.44`, exceeds the maximum possible surplus contribution `eta ln N = 0.8697`.

Steady states, for completeness:

```
Q_inf = (eta * F(u) - C) / delta
at u = 0.28:  (0.35*1.406097 - 0.40) / 0.08 = 1.1517
at u = 0.41:  (0.35*1.706565 - 0.40) / 0.08 = 2.4662
at u = 0.52:  (0.35*1.905088 - 0.40) / 0.08 = 3.3348
```

The institute is carrying `Q = 18.0` and its best achievable steady state at `u = 0.52` is 3.33. It is not underperforming; it is over-extended. The intervention that raises the tone by a minor seventh does not fix this and cannot. Chapter 6's four levers apply: raise `eta`, raise `N`, cut `C`, or **shed `Q`** — and here the arithmetic points hard at the last, which is the lever nobody considers.

This is precisely the failure FM-30.9 names. Had the institute adopted the cross-disciplinary tone as its KPI, it would have celebrated a 1040-cent gain through the two years in which its stored order was decaying toward a fixed point one-fifth of its current size.

Step 9 — Falsification (Chapter 25). The prediction, preregistered before the rotation programme began: u measured by the frozen protocol will reach `0.41 +/- 0.06` within eight quarters, and the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) − C_t + eps` fitted on quarters 1-8 and frozen will beat all four required baselines (persistence, direct-input, matched autoregressive, and the shuffled-input control) on quarters 9-12. **If the shuffled-input control performs as well as the fitted model, the surplus term is carrying no information and the tone reported above is an arithmetic exercise on noise. That is the outcome that would falsify this worked example, and it is stated before the data are in.**

Step 10 — Clauses (BR-2, BR-3, BR-10), stated in full as required.

The identification `n_eff = N_eff` is a modeling decision, warranted by the six-requirement argument of §30.4 and not derived. The affine effective-count axiom S3 on which `N_eff` depends is a design axiom, chosen rather than derived. Every frequency above is an essence frequency, not a physical resonance: no claim is made that the institute, its researchers, its rooms or its disciplines physically resonate at any of these values. The surplus functional as deployed is a lower bound on the ensemble's information, so every tone quoted here is a lower bound on the pitch. `N = 12` is one partition among many and is not asserted as correct; changing it changes every number in this example.

Chapter 31 takes these tones and builds with them: a room whose dimensions come out of an essence set, a cadence whose period comes out of the same arithmetic, and a specification — stated there in the strongest terms this manual uses — that everything sounded is **audible and disclosed**, and that nothing is ever masked or hidden from a listener.

Chapter 31 . Sonic Alchemy in Built Space

Chapter 29 gave the essence law. Chapter 30 gave the bridge that lets it cross a discipline boundary. This chapter builds with them.

The claim being made is narrow and should be read narrowly. **The essence law is a target-generating instrument.** It takes an ordered category scheme — a room's programme, an institution's houses, a working group's roles — and returns a small set of numbers with exact ratio relationships. Those numbers are then used the way any designer uses a proportional system: to choose dimensions, to choose tuning targets, to choose a cadence. The physics that follows is ordinary physics. Sabine's equation is Sabine's equation. A room 4.38 metres wide has an axial mode at 39.16 Hz because `f = c/(2L)`, not because of anything in this book.

What the essence law contributes is the choice of target, and nothing else. It replaces "pick a proportion that looks right" with "pick the proportion your own category scheme generates, and be able to state where it came from". That is a real contribution and it is a modest one, and this chapter will not claim more.

The chapter carries one constraint that outranks every design consideration in it, stated first because it must never be reached as an afterthought.

The Law

31.1 The no-subliminals constraint — binding, absolute, and prior to everything else

No subliminal content of any kind. Nothing masked. Nothing below the threshold of hearing. Nothing concealed from the listener. A message hidden from its recipient cannot be consented to.

That is the house design constraint (Methodology Protocol §8.1), and in this chapter it is the first law rather than a caveat, because this is the chapter where the temptation arises. Every technique in this chapter operates **entirely in audible, disclosed tone.**

Concretely, and as a specification rather than a sentiment:

1. **Every tone present in a space must be audible to a person of ordinary hearing standing in it.** If a tone is below the threshold of audibility at any listening position, it is either raised to audibility or removed. There is no third option.
2. **Every tone must be disclosed.** A written statement — posted at the entrance, or in the room's documentation, or both — lists what is sounding, at what frequency, at what level, and why. A person must be able to know what is in the room before deciding to be in it.
3. **No masking, no carrier, no embedding.** No content is placed under, inside, or behind another signal. No amplitude-modulated speech, no time-compressed content, no spectral hiding, no content below 20 Hz or above 20 kHz intended to act on a listener.
4. **The room must be leavable and the tone must be stoppable.** There is an off switch, its location is disclosed, and anyone in the room may use it. A tuned space that a person cannot switch off or walk out of is a coercive instrument, and CANON §5.4 puts it out of scope for this manual.

5. **Where the tuning is structural rather than sounded — a room dimension, a meeting cadence — the derivation is published.** A proportion chosen by a method should be attributable to that method in writing, so that anyone can disagree with it.

The 2022 source document recommends embedding ASMR subliminal messages in the rendered track. **This chapter does not carry that recommendation and specifically excludes it.** The reason is not squeamishness and not doubt about efficacy — efficacy is beside the point. **A method that works by removing a person's ability to decline is out of scope for this manual regardless of whether it works, and a message the listener cannot hear is a message the listener cannot decline.** Chapter 20 makes the same argument structurally: the house method is to build a lower valley next door, never to push. Its worst outcome is that someone declines to move, and that is a safety property, not a weakness.

The same source recommends a 12- or 24-winding Rodin coil as an amplifier, on the stated grounds that it "influences temporal 4th dimensional space by increasing the probability of what you play through the coil over time". **This chapter does not present the Rodin coil as an established amplifier.** There is no verified engineering basis for the claimed mechanism, and CANON §5.3 forbids offering a number or a device that has not been computed or cited. If you want an area of effect, use a loudspeaker with a published frequency response and a measured sound pressure level, and disclose both.

31.2 The three laws this chapter builds on

POSTULATE (Essence Law, Chapter 29). For an essence index n in a defensibly ordered category set,

```
f(n) = k * n^2,   k = (9/8)^2 / phi^2 = 0.483425732988414 Hz    [badged: k is chosen]
```

PROVED (Result One). `f(m)/f(n) = m^2/n^2`. Every interval is a ratio of squares and k cancels.

PROVED (Result Two). Two essences share a pitch class iff their indices have the same odd part; a set indexed `1..N` has exactly `ceil(N/2)` distinct pitch classes.

IDENTIFICATION (Chapter 30, badged). Across disciplines, `n_eff = N_eff = 1 + (N−1)u = e^F`, so `f_cross = k e^(2F)` and `octaves above base = 2F/ln 2 = 2.885390081777927 * F`.

PROVED (Sabine). `RT60 = 0.161 * V / A` (metric), V in cubic metres, A in square-metre sabins. Real acoustics, Chapter 14.

PROVED (room modes). For a rigid rectangular room, `f(p,q,r) = (c/2) * sqrt((p/Lx)^2 + (q/Ly)^2 + (r/Lz)^2)`, `c = 343 m/s` at 20 C in dry air. Real acoustics.

31.3 The honesty clauses, restated in the built context

Chapter 31 is where the clauses are easiest to forget and most expensive to forget, because here the numbers become concrete objects.

H1 — essence, not resonance. The output of the essence law is an essence frequency. **No claim is made that any material, room, institution or person physically resonates at these values.** A room whose axial mode is placed at 39.157484 Hz genuinely resonates at 39.157484 Hz — that is ordinary acoustics, and it is true because the room is 4.3798 metres wide. It is *not*

true because the room's essence index is 9. **The essence law chose the target; physics did the resonating; those are two different sentences and the manual keeps them apart.**

H2 — k is chosen. Badged everywhere.

H4 — the ordering must be defensible. The whole method is downstream of §29.9's test. Run it.

H5 — recompute; never copy a published value. This clause matters more here than anywhere else in Part V, because a built object freezes an arithmetic error in timber and plaster and nobody re-checks it afterwards. §29.4 established that `f = k n^2` reproduces all 118 published element values to a worst relative deviation of **9.0e-15**, and §29.10 that the compound rule `f = k^2 G^4 — G` the count-weighted geometric mean of the constituent indices — matches **37 of the 44** published composites to better than one part in 10^12. **The seven that disagree are data-entry errors in the source, not a different law, and the O2 row proves it: for a single-element species `G` is the element index whatever the atom count, so no atom count whatever can produce the published figure.** Two transcription faults travel with any citation of the source: its prose instruction `1+5^.5/2` is missing its parentheses and must be read as `(1 + sqrt(5))/2`, and the spreadsheet stored `phi` to fifteen significant figures. **Cut a room, a bar or a pipe to a recomputed value and cite the errata beside it.**

From the surplus framework. S3 is a design axiom. Surplus as deployed is a lower bound. No thermodynamic claim follows. `N` must be published with every value.

And one clause specific to building: *the fact that a room is tuned by this method is not evidence that the method works.* A room built to these proportions may be a good room for entirely conventional reasons — its ratio happened to land near Volkmann's, its RT60 happened to land in the speech range, its size happened to suit twelve people. **Chapter 25's falsification discipline applies: if you want to claim the essence tuning did something, you need a preregistered prediction and the four required baselines, including the shuffled-input control. This chapter claims only that the method produces a defensible, reproducible, documented choice of target.** And carry the correction with the requirement: Chapter 25 §25.8 records three measured failure modes of the shuffled-input control — it is anti-conservative as ordinarily run, it cannot see repackaging, and it cannot see a shared slow driver — so a passed shuffle is not on its own evidence that a tuned room did anything. **Design controls, not more surrogates.**

The Engineering

31.4 The procedure

Six steps. Every one of them is checkable by a third party from the published record.

```
  STEP 1  Declare the category set and the ordinal rule.      (Paper B 1.2; ES-
3)
   STEP 2   Run the permutation sensitivity test.                 (ES-4)
   STEP 3   Compute raw essence frequencies f(n) = k n^2.
   STEP 4   Reduce to pitch classes by odd part; count ceil(N/2). (ES-6)
   STEP 5   Fold into the working band; publish band and j.       (ES-7)
   STEP 6   Apply to one or more of the four targets below.
```

The four targets, each treated in its own section:

- **(a)** acoustic tuning targets, interacting with Chapter 14's Sabine calculations — §31.5
- **(b)** room dimension ratios and modal frequencies — §31.6
- **(c)** scheduling and cadence, as Chapter 2's periodic forcing of a low-pass filter — §31.7
- **(d)** tuned physical elements — §31.8

Step 5 deserves a note, because it is where most of the practical judgement lives. Raw essence frequencies span from sub-audible to ultrasonic: for a twelve-set they run 0.483426 Hz to 69.613306 Hz, entirely below the useful speech band. Octave transposition — multiplying by `2^j` — preserves pitch class exactly and is the house method. The genomic-audio work in this lineage uses a 36-octave transposition on the same principle. **State the band and the `j` for every value published; a folded value without its `j` is not comparable to anything (FM-29.5).**

The house band for structural work in this chapter is `[220, 440)` Hz — one octave, anchored on A3 = 220 Hz. It was chosen because it sits inside the speech fundamental range for adult voices, is well above typical room-mode territory, is comfortably below the 2-4 kHz consonant band where intelligibility lives, and folds a twelve-set into a single octave with `j` between +2 and +9. **That is a design choice, stated as one.**

31.5 Target (a): acoustic tuning targets and the Sabine interaction

Chapter 14 established the framework: `RT60 = 0.161 V / A` in metric units, Woolley's collective intelligence factor tracks equality of conversational turn-taking, and turn-taking is substantially a property of the room. **A hard-surfaced room silences the quiet voice by physics, not by dominance.** Acoustic design is c-factor engineering and it is the cheapest such intervention available.

The essence law adds one thing to that: it tells you **which octave bands to treat first**.

Absorption is not flat. Porous absorbers (mineral wool, heavy curtain, carpet) are efficient above about 500 Hz and become progressively useless below 250 Hz; membrane and Helmholtz absorbers are narrowband and must be tuned. So the practical question in any real treatment budget is *which bands do I spend on?* — and the essence set answers it by identifying where the room's own programme places its pitch classes.

Worked, for a twelve-set folded into `[220, 440)`:

class	folded Hz	octave band	inside the 250-4000 Hz speech range?
PC1	247.5140	250	at the lower band edge
PC3	278.4532	250	yes
PC5	386.7406	500	yes
PC7	379.0058	500	yes
PC9	313.2599	250	yes
PC11	233.9781	250	at the lower band edge

Four of six classes land in the 250 Hz octave band and two in the 500 Hz band. That is

the treatment priority, and it is a useful one: the 250 Hz band is exactly where cheap porous absorption starts losing efficiency and where most under-treated meeting rooms have their worst speech-intelligibility problem. The essence set is pointing at the band that most rooms neglect.

The Sabine arithmetic for a 100 m^3 room, computed:

A (m^2 sabins)	RT60 (s)
8	2.013
10	1.610
13	1.238
16	1.006
20	0.805
26	0.619
32	0.503
40	0.403

and inverted, the absorption required to hit a target:

```
V = 100 m^3
RT60 = 0.4 s  ->  A = 40.25 m^2 sabins
RT60 = 0.5 s  ->  A = 32.20 m^2 sabins
RT60 = 0.6 s  ->  A = 26.83 m^2 sabins
RT60 = 0.8 s  ->  A = 20.12 m^2 sabins
RT60 = 1.0 s  ->  A = 16.10 m^2 sabins
```

One hard limit that the essence law cannot touch and must not be allowed to obscure. The Schroeder frequency,

```
f_s = 2000 * sqrt(RT60 / V)
```

marks the transition between modal behaviour (discrete resonances, position-dependent response) and statistical behaviour (diffuse field, where Sabine applies). Computed for `V = 100 m^3`:

```
RT60 = 0.40 s  ->  f_s = 126.5 Hz
RT60 = 0.50 s  ->  f_s = 141.4 Hz
RT60 = 0.60 s  ->  f_s = 154.9 Hz
RT60 = 0.80 s  ->  f_s = 178.9 Hz
```

Below `f_s`, Sabine's equation does not apply and the room's behaviour is set by individual modes. Above it, modes are dense and Sabine governs. Since the folded twelve-set lives between 234 and 387 Hz, it sits entirely in the statistical region for a room of this size — which is the correct place for it, and which is why the `[220, 440)` band was chosen. **A folded band that straddled `f_s` would produce a set of targets half of which obeyed one physics and half another, and the resulting specification would be incoherent.** Check `f_s` before choosing the band. Always.

31.6 Target (b): room dimension ratios and modal frequencies

This is the strongest of the four applications, because room proportion is a genuinely open design question with a wide space of acceptable answers, and the essence law narrows it defensibly.

An axial mode along a dimension L sits at `f = c/(2L)`, so `L = c/(2f)`. Choose three essence classes, one per axis, and the room's three axial fundamentals *are* those essence frequencies. The ratio of the dimensions is then the inverse ratio of the class frequencies — a ratio of squares by Result One, hence always a rational number with an exact expression.

The constraint that makes this non-trivial is that most ratios are bad rooms. Room-ratio heuristics exist because degenerate modes — two or three modes at the same frequency — create strong peaks and long decays at those frequencies, and near-degenerate axial modes create audible colouration of speech. The published heuristics (they are heuristics, and are stated here as such) include:

```
Louden          1 : 1.4  : 1.9
Sepmeyer A      1 : 1.14 : 1.39
Sepmeyer B      1 : 1.28 : 1.54
Sepmeyer C      1 : 1.60 : 2.33
Volkmann        1 : 1.5  : 2.5
```

So the design question becomes: **is there an assignment of essence classes to axes whose ratio lands near a recommended ratio?** This was searched exhaustively over all class triples and all octave offsets from 0 to 6, for the six classes of a twelve-set. Best matches, computed:

heuristic target	nearest essence ratio	classes used	Euclidean error
Louden 1 : 1.40 : 1.90	1 : 1.3611 : 1.9600	PC3, PC5, PC7	0.0715
Sepmeyer A 1 : 1.14 : 1.39	1 : 1.1250 : 1.4400	PC1, PC3, PC5	0.0522
Sepmeyer B 1 : 1.28 : 1.54	1 : 1.2346 : 1.5625	PC1, PC5, PC9	0.0507
Sepmeyer C 1 : 1.60 : 2.33	1 : 1.6200 : 2.2500	PC3, PC5, PC9	0.0825
Volkmann 1 : 1.50 : 2.50	**1 : 1.4938 : 2.4694**	**PC7, PC9, PC11**	**0.0312**

The best match is PC7 : PC9 : PC11 against Volkmann, with a Euclidean error of 0.031 — the middle term within 0.4% and the long term within 1.2% of the heuristic. And the ratio has an exact closed form, which is the whole point of using the essence law rather than a table lookup:

```
L proportional to 1/f, and f proportional to m^2, so

L(PC11) : L(PC9) : L(PC7)  =  1/121 : 1/81 : 1/49
                           =  1 : 121/81 : 121/49
                           =  1 : 1.493827... : 2.469388...
```

The ratio is exactly `1 : 121/81 : 121/49`. Not approximately. That is what Result One buys: **every proportion this method produces is a ratio of squares of integers and can be written**

down exactly, which means it can be checked, argued with, and reproduced by someone who does not trust you.

Now anchor the absolute size by requiring the axial fundamentals to *be* the essence frequencies:

```
f(7)  = k * 49  = 23.687861 Hz  ->  Lz = 343/(2*23.687861) = 7.2400 m   (length)
f(9)  = k * 81  = 39.157484 Hz  ->  Ly = 343/(2*39.157484) = 4.3798 m   (width)
f(11) = k * 121 = 58.494514 Hz  ->  Lx = 343/(2*58.494514) = 2.9319 m   (height)

V = 92.969 m^3      floor area = 31.709 m^2      total surface S = 131.555 m^2
  occupancy at 2.5 m^2 per person = 12.7 people
```

A room 2.93 m high, 4.38 m wide and 7.24 m long, holding about twelve people. The height is buildable, the floor plan is a conventional meeting-room shape, and the occupancy matches the twelve-house set it was derived from. None of that was forced; it fell out of `c/(2f)` with `c` = 343 m/s. It is a genuinely lucky landing and it should be said plainly that it is lucky — a different k or a different class triple would have produced an unbuildable room, and §31.11's failure modes cover what to do then.

The modal structure, computed for this room, first twenty modes:

#	f (Hz)	(p,q,r)	type	gap from previous (Hz)	essence relation
1	23.69	(0,0,1)	axial	-	`f(7) x 2^0`
2	39.16	(0,1,0)	axial	15.47	`f(9) x 2^0`
3	45.76	(0,1,1)	tangential	6.61	-
4	47.38	(0,0,2)	axial	1.61	`f(7) x 2^1`
5	58.49	(1,0,0)	axial	11.12	`f(11) x 2^0`
6	61.46	(0,1,2)	tangential	2.97	-
7	63.11	(1,0,1)	tangential	1.65	-
8	70.39	(1,1,0)	tangential	7.28	-
9	71.06	(0,0,3)	axial	0.67	-
10	74.27	(1,1,1)	oblique	3.21	-
11	75.27	(1,0,2)	tangential	1.00	-
12	78.31	(0,2,0)	axial	3.04	`f(9) x 2^1`
13	81.14	(0,1,3)	tangential	2.82	-
14	81.82	(0,2,1)	tangential	0.68	-
15	84.85	(1,1,2)	oblique	3.03	-
16	91.53	(0,2,2)	tangential	6.68	-
17	92.04	(1,0,3)	tangential	0.51	-
18	94.75	(0,0,4)	axial	2.71	`f(7) x 2^2`
19	97.75	(1,2,0)	tangential	3.00	-
20	100.02	(1,1,3)	oblique	2.28	-

```
modes below 150 Hz:              48
mean gap of the first 20 modes:  4.02 Hz
maximum gap:                     15.47 Hz  (between modes 1 and 2)
minimum gap:                     0.51 Hz   (modes 16 -> 17)
```

Read the "essence relation" column and Result Two is visible in built form. **The room's axial mode series along each axis is `f, 2f, 3f, 4f, ...`, and the even members of each series are the essence frequencies of that axis's class at higher octaves.** Mode 4 is `f(7) x 2`, mode 18 is `f(7) x 4`, mode 12 is `f(9) x 2`. The odd members — mode 9 at 71.06 Hz, which is `3 x f(7)` — are *not* essence frequencies of anything in the set, because tripling an index is not an operation the essence law preserves. **The room delivers the essence set at the even harmonics and delivers ordinary harmonic series members between them. Do not report the odd members as essence frequencies; they are not.**

Three honest acoustic observations about this specific room, none of which the essence law predicted and all of which conventional analysis produces:

1. **The 15.47 Hz gap between modes 1 and 2 is large.** A single isolated mode at 23.69 Hz with nothing near it will ring. At this frequency it is below the useful speech range and mostly inaudible as pitch, but it will couple to footfall and HVAC. Treat with a membrane absorber or accept it.
2. **The 0.51 Hz gap between modes 16 and 17 is a near-degeneracy.** Two modes at 91.53 and 92.04 Hz will beat at about half a hertz — a slow throb. Both are tangential, so the effect is mild, but it is real and it is worth knowing before someone reports the room as "wobbly" and nobody knows why.
3. **48 modes below 150 Hz** is a healthy density for a room of this volume; the low end is not sparse. The Schroeder frequency at `RT60 = 0.5 s` is `2000 * sqrt(0.5/92.969) = 146.7 Hz`, consistent with that count.

Sabine treatment for this room, computed:

target `RT60` (s)	required `A` (m^2 sabins)	mean absorption `alpha_bar = A/S`
0.4	37.42	0.2844
0.5	29.94	0.2276
0.6	24.95	0.1896
0.7	21.38	0.1625
0.8	18.71	0.1422
1.0	14.97	0.1138

Chapter 14's deliberation-room target is met at `RT60 = 0.5–0.6 s`, requiring 25 to 30 square-metre sabins across a 131.6 m^2 surface — a mean absorption coefficient of 0.19 to 0.23. That is achievable with a carpeted floor, an absorbent ceiling raft over the table, and one treated end wall; it does not require a specialist build. The specification in §31.10 states it.

31.7 Target (c): scheduling and cadence

Chapter 2 established that the storage law is a first-order low-pass filter with pole at `(1–delta)` and time constant `tau` of about `1/delta` periods, and that **ritual is a periodic forcing function on a low-pass filter, which is why cadence beats intensity.** The essence law can set the period.

The mechanism is octave folding applied to time rather than to frequency. An essence frequency `f` has a period `T = 1/f` seconds. Multiply by `2^p` — the same operation as octave transposition,

since period is the reciprocal of frequency — to bring T into a humanly usable band. Pitch class is preserved exactly, which here means: **two essences that share a pitch class share a cadence, up to a doubling.** That is a defensible property for a scheduling system, because two categories that are harmonically identical should not be given unrelated meeting rhythms.

Folded into `[1 day, 2 days)`:

class	f raw (Hz)	T = 1/f (s)	p	T ∗ 2^p (s)	days
PC1	0.483426	2.068570	16	135565.8	1.5690
PC3	4.350832	0.229841	19	120502.9	1.3947
PC5	12.085643	0.082743	20	86762.1	1.0042
PC7	23.687861	0.042216	21	88532.8	1.0247
PC9	39.157484	0.025538	22	107113.7	1.2397
PC11	58.494514	0.017096	23	143408.5	1.6598

Folded into `[7 days, 14 days)` — the more useful band for institutional cadence:

class	p	T ∗ 2^p (days)	practical reading
PC5	23	8.0335	weekly, drifting one day per cycle
PC7	24	8.1975	weekly, drifting one day per cycle
PC9	25	9.9179	ten-day
PC3	22	11.1577	eleven-day
PC1	19	12.5524	fortnightly, slightly early
PC11	26	13.2786	fortnightly, slightly early

Two things to say about this table, one useful and one honest.

The useful thing: **the spread from 8.03 to 13.28 days is a factor of 1.65, and a set of cadences in that spread will drift in and out of phase with each other on a cycle of months.** That is a design property, not a defect. Two committees meeting every 8.03 and 9.92 days coincide every `8.0335 ∗ 9.9179 / (9.9179 − 8.0335) = 42.3` days — roughly every six weeks — which is a natural rhythm for a joint review. Chapter 21's stigmergy argument applies: **you can schedule a coincidence without a coordinator, by choosing periods whose beat frequency is the interval you want.**

The honest thing: **nothing here is calendar-compatible, and calendars are strong.** A 9.92-day cycle walks through the week and will land on a Sunday about one time in seven. In practice you round to the nearest workable integer — 8, 10, 11, 14 days, or "every second Tuesday" — and when you do, you have left the essence law behind and are simply scheduling. **Say so. A rounded cadence is a rounded cadence and claiming it is essence-derived after rounding is FM-31.5.**

The relation to Chapter 2's filter is where this becomes engineering rather than numerology. What matters is the ratio of the forcing period `P` to the system's time constant `tau = 1/delta`:

delta (per week)	tau (weeks)	settling to 95% ~ 3 tau (weeks)	P/tau at P = 7.29 weeks
0.02	50.00	150.00	0.146
0.05	20.00	60.00	0.364
0.10	10.00	30.00	0.729
0.20	5.00	15.00	1.457
0.30	3.33	10.00	2.186

The rule that comes out of this is the useful one, and it does not depend on the essence law at all: choose P well below tau, ideally P/tau <= 0.5. A low-pass filter driven well below its corner frequency integrates smoothly and holds a near-constant Q; driven near or above it, Q sawtooths — building after each ritual and decaying visibly between them, which is the felt experience of an institution whose practices "don't stick". At delta = 0.10 per week, a 51-day (7.29-week) cadence gives P/tau = 0.729, which is marginal. At delta = 0.05 the same cadence gives 0.364, which is comfortable.

The essence law proposes the cadence; delta decides whether it is admissible. If the essence derived period exceeds tau/2, fold it down one more octave and use the half-period. That is a legal move — it preserves pitch class exactly — and it is the correct one. Halving the PC1 cadence from 12.55 days to 6.28 days is not an abandonment of the method; it is the method's own octave operation applied to time.

For completeness: the 2022 source specifies a tempo rule — multiply the essence frequency by 60 and halve until the result falls between 100 and 200, then use it as the track's BPM. Recomputed:

class	f raw (Hz)	f * 60	halvings	BPM
PC1	0.483426	29.0055	-2	116.0222
PC3	4.350832	261.0499	+1	130.5249
PC5	12.085643	725.1386	+2	181.2846
PC7	23.687861	1421.2717	+3	177.6590
PC9	39.157484	2349.4491	+4	146.8406
PC11	58.494514	3509.6708	+5	109.6772

This is octave folding of the period into a musically conventional band, which is exactly the same operation as the day-band folding above with a different target window. It is recorded because it is the source's own procedure and because the numbers should be available; it carries no additional claim.

31.8 Target (d): tuned physical elements

Physical resonators tuned to essence classes give a room a set of objects that sound its own proportional system when struck, blown, or excited by the room itself. They are also the cheapest way to make the tuning *disclosable*: a visible object with a label on it satisfies §31.1's disclosure requirement in a way a hidden loudspeaker never can.

All three formulas below are ordinary physics and none of them is contributed by the essence law. The essence law supplies only the target frequency.

Open-open tubes (both ends open), `f = c / (2*(L + 0.8*d))` with `d` the bore diameter, for a 50 mm bore, `c = 343` m/s:

class	target `f` folded to [220,440) (Hz)	geometric length `L` (m)
PC1	247.5140	0.6529
PC3	278.4532	0.5759
PC5	386.7406	0.4034
PC7	379.0058	0.4125
PC9	313.2599	0.5075
PC11	233.9781	0.6930

Helmholtz resonators, `f = (c/(2*pi)) * sqrt(A / (V * (L_neck + 1.7*a)))`, neck radius `a = 0.025 m`, neck length 0.10 m, targeting one octave below the folded band (`f/4`) so the units are practical bass-trap sizes:

class	target `f/4` (Hz)	required cavity `V` (m^3)	litres
PC1	61.8785	0.010724	10.72
PC3	69.6133	0.008473	8.47
PC5	96.6851	0.004393	4.39
PC7	94.7514	0.004574	4.57
PC9	78.3150	0.006695	6.70
PC11	58.4945	0.012001	12.00

Note that the PC11 target, 58.4945 Hz, is exactly `f(11)` unfolded — the folded-band value happened to be four octaves above it. And the PC3 target, 69.6133 Hz, is exactly `f(12) = k * 144`. **Both coincidences are Result Two doing its work: 11 and 44 share an odd part, as do 3 and 12.** A Helmholtz trap built to 58.49 Hz sits within 0.005 Hz of the room's PC11 axial mode at 58.494514 Hz (§31.6), which makes it a genuinely useful modal trap and not merely a decorative one.

Stretched strings / monochord, `f = (1/(2L)) * sqrt(T/mu)`, steel with `mu = 0.0012 kg/m` at `T = 400 N`, so the wave speed is `sqrt(400/0.0012) = 577.35 m/s`:

class	target `f` (Hz)	speaking length `L = v/(2f)` (m)
PC1	247.5140	1.1663
PC3	278.4532	1.0367
PC5	386.7406	0.7464
PC7	379.0058	0.7617
PC9	313.2599	0.9215
PC11	233.9781	1.2338

A six-string monochord with speaking lengths from 0.746 to 1.234 metres is a buildable object about 1.3 metres long — a bench-scale instrument that can sit along a wall, be struck by anyone,

and carry a printed label giving all six frequencies, their essence indices, and the ordinal rule they came from. **That single object satisfies §31.1 clauses 1, 2 and 5 simultaneously: it is audible, it is disclosed, and its derivation is published on it.** It is the recommended physical form for this method, in preference to any loudspeaker system, for exactly that reason.

31.9 Chords, beats, and what the tuning actually sounds like

Sound the six folded classes together and this is what happens, computed:

```
sorted folded frequencies (Hz):
233.9781   247.5140   278.4532   313.2599   379.0058   386.7406
```

Pairwise differences — the beat rates — for the two closest pairs:

```
379.0058 & 386.7406  ->  difference  7.7348 Hz   ROUGHNESS ZONE
233.9781 & 247.5140  ->  difference 13.5359 Hz   ROUGHNESS ZONE
all other pairs      ->  30.9 to 152.8 Hz, heard as separate pitches
```

Two of the fifteen pairs beat inside the roughness zone (difference under about 20 Hz), and PC7 against PC5 at 7.73 Hz is a distinctly audible flutter rather than a chord. This is a direct consequence of Chapter 29's §29.8: PC7's folded ratio is 1.531250 and PC5's is 1.562500, a gap of 34.98 cents — a third of a semitone. The two classes are structurally almost-identical and the ear hears that as beating rather than as harmony.

This is a finding about the category scheme, not a defect in the method. It says something true: in a twelve-house partition, the houses at indices 5 and 7 (and 10 and 14, and so on) sit at nearly the same harmonic position. If you want them to be distinguishable, you sound them separately, or you transpose one by a full octave so the 34.98-cent gap becomes a 1234.98-cent interval and reads as an octave-plus rather than a beat. **Both are legal moves. Concealing the beat by dropping one voice below audibility is not (§31.1).**

The three usable simultaneities in this set, in order of consonance:

```
   PC1 + PC3    ratio 9/8,      203.91 cents  Pythagorean whole tone, +3.91 from 12-
TET
   PC1 + PC9    ratio 81/64,     407.82 cents  Pythagorean ditone, +7.82 from 12-
TET
   PC1 + PC11   ratio 121/64,   1102.64 cents  near-just major seventh, +2.64 from 12-
TET
```

All three are within 8 cents of a twelve-tone equal-tempered interval, which means this set will sit comfortably against conventional instruments without retuning them. That is a practical convenience and it is a consequence of 9/8 and 121/64 being close to tempered intervals; it is not a claim about anything.

Three rendering constraints, from Chapter 71's Gabor laws, that bind anything actually built.

The set size is capped, and twelve is safely inside the cap. Under a single integer octave transposition into 20 Hz - 20 kHz, at most **25 essences** fit, and only at `j = 6` (Chapter 29 §29.8, Law 1). A twelve-house set fits with room to spare. A thirty-two-system set does not, and must be rendered under split transposition — which changes the interval relations between the two halves and must be declared, not quietly applied.

Do not grow the set past sixteen without expecting the flutter above to get worse. The 34.98-cent gap between PC7 and PC5 is the *widest* minimum gap available: it holds for every set size from 7 to 16. **The moment a seventeenth index enters, the minimum gap collapses to 6.001 cents** — index 17 folds to 289/256, six cents from index 3's 9/8 — and against PC3 at 278.4532 Hz that is a beat of **0.97 Hz, against the present 7.73 Hz**, which is no longer a flutter but a slow swell that most listeners will hear as one wandering tone. Both of the remedies above still apply, and both cost more at a seventeen-house scale than at twelve.

Every tone needs a minimum duration, and a rectangular gate is not a legal envelope. Resolving a surplus difference `dF` requires at least `1/(8 pi dF)` cycles (Law 3), independent of transposition: separating PC7 from PC5 needs 3.94 cycles, about 10 ms at 379-387 Hz, while separating a 6.001-cent pair at PC3's 278.4532 Hz would need 22.96 cycles, about 82 ms. **Specify the tone duration as a cycle count, not as a wall-clock figure**, because the cycle count is what transposes. And the envelope itself is constrained: a rectangular on/off gate has **no finite spectral second moment** — its `sinc^2` tails fall as `1/f^2` and the second-moment integral is numerically divergent — so it has no defined bandwidth and **is disqualified as a rendering envelope** (Law 4). Use a windowed envelope: a raised-cosine rise and fall of at least one cycle at each end is sufficient and is the house default. This is an engineering requirement, not an aesthetic preference, and it is entirely compatible with §31.1 — a windowed envelope is audible, disclosed, and hides nothing.

The Specification

SA-1 — NO SUBLIMINALS. No masked content, no content below the threshold of hearing at any listening position, no content concealed from a listener, no carrier or embedding of any kind. This requirement outranks every other requirement in this chapter and admits no exception. (§31.1)

SA-2 — Disclose in writing. A posted statement lists every tone present, its frequency, its measured level in dB SPL at the reference position, and its derivation. Posted at the entrance, readable before entry.

SA-3 — Provide an off switch and disclose its location. Anyone in the space may stop the tone. A tuned space that cannot be switched off or walked out of is out of scope for this manual.

SA-4 — Do not present the Rodin coil as an amplifier. No verified engineering basis exists for the claimed mechanism. Use a loudspeaker with a published frequency response and a measured SPL, and disclose both. (§31.1)

SA-5 — Publish the ordinal rule and the permutation sensitivity result before publishing any frequency. (ES-3, ES-4)

SA-6 — Publish `(N, class, raw Hz, j, band, folded Hz)` for every tone. A folded value without its `j` is not comparable to anything. (ES-7)

SA-7 — Design to `ceil(N/2)`, not to `N`. Twelve houses give six pitch classes. (ES-6)

SA-8 — Check the Schroeder frequency before choosing the band. `f_s = 2000 sqrt(RT60/V)`. The whole folded set must sit on one side of `f_s`, and for speech work it must sit above it. (§31.5)

SA-9 — Compute the full modal set, not just the axials. Report modes below 150 Hz, the maximum gap, and the minimum gap. Flag any gap below 1 Hz as a near-degeneracy with its predicted beat rate. (§31.6)

SA-10 — Meet Chapter 14's RT60 target first; use the essence set only to prioritise bands. If the essence-derived proportion and the RT60 target conflict, **RT60 wins.** Speech intelligibility is a measured requirement; proportion is a design choice.

SA-11 — Verify that essence-derived cadences satisfy `P <= tau/2`, `tau = 1/delta`. If not, fold down one octave and use the half-period. (§31.7)

SA-12 — State every rounding. A cadence rounded to a calendar, a dimension rounded to a construction module, a frequency rounded to a tuner's resolution — each is a departure from the method and each is disclosed as one.

SA-13 — Prefer a labelled physical resonator to a loudspeaker. A struck object with its derivation printed on it satisfies audibility, disclosure and reproducibility at once. (§31.8)

SA-14 — Report beat rates for every pair under 20 Hz apart. Roughness is a real perceptual effect and a room that flutters must say that it flutters. (§31.9)

SA-15 — Restate H1 on every drawing, schedule and label. "Essence frequency, not physical resonance. No claim is made that any material, room, institution or person physically resonates at these values."

SA-16 — Preregister before claiming an effect. If you intend to claim the tuning changed anything, Chapter 25's protocol applies in full, including the shuffled-input control. Absent that, the only claim available is that the target was chosen by a stated, reproducible method. (§31.3)

The Failure Modes

FM-31.1 THE HIDDEN TONE

Damages: consent, and with it the entire ethical standing of the work. *Description.* A tone, a message, or a modulation is placed below the threshold of audibility, or masked under a louder signal, or embedded in a carrier. *Signature.* Any signal path that a person in the room cannot hear and was not told about; any spectrum analyser trace showing content not on the posted disclosure. *Fix.* SA-1, SA-2, SA-3. **There is no partial fix and no acceptable dose. Remove it or make it audible and disclosed.**

FM-31.2 THE RESONANCE CLAIM, BUILT

Damages: everything, and here it is easiest to commit because the room really does resonate. *Description.* "The room resonates at the essence frequency of Biology & Body." The room resonates at 39.157484 Hz because it is 4.3798 m wide. Biology does not resonate at anything. *Signature.* An essence, a discipline, a material or a person as the grammatical subject of "resonates". *Fix.* SA-15, and rewrite: "the width was chosen so that the axial mode falls at the essence frequency of index 9."

FM-31.3 THE UNBUILDABLE ROOM

Damages: the project. *Description.* A class triple is chosen for harmonic reasons and produces dimensions of, say, `0.74 x 3.60 x 10.00 m` — a ceiling height of 74 centimetres and an aspect ratio of 13.4 : 1. Real: that is exactly what PC11 : PC5 : PC3 gives if the raw class ratios are used with no octave offset, scaled so the middle axis is 3.60 m. The ratio `1 : 121/25 : 121/9` is harmonically legitimate and architecturally impossible. *Signature.* Any dimension below about 2.4 m or above about 5 m for a room-scale build; an aspect ratio beyond about 1 : 3. *Fix.* Search the class triples *and* the octave offsets, as §31.6 does, before committing. If no combination is buildable, the honest report is that this category scheme does not produce a room, and the essence set is used for cadence or tuned elements instead.

FM-31.4 THE MODAL BLIND SPOT

Damages: the acoustic outcome, badly. *Description.* Axial modes are placed by essence and the tangential and oblique modes are never computed. Two of them land 0.5 Hz apart and the room throbs. *Signature.* A room-mode analysis containing only `(p,0,0)`, `(0,q,0)`, `(0,0,r)` entries. *Fix.* SA-9. Compute the full set to at least `p,q,r = 4` and report the gaps.

FM-31.5 THE ROUNDED CADENCE, UNDISCLOSED

Damages: the reproducibility of the whole schedule. *Description.* An essence-derived 9.9179-day cadence is rounded to "every second Tuesday" (14 days) and the schedule is still described as essence-derived. It is not; 14 days is 1.41 times the derived value, a difference of 590 cents. *Signature.* A published cadence at a round calendar interval with an essence derivation attached. *Fix.* SA-12. State the derived value, state the rounding, state the size of the departure. A rounded cadence is a perfectly good cadence; it is just not the derived one.

FM-31.6 PROPORTION OVER INTELLIGIBILITY

Damages: speech, turn-taking, and therefore c (Woolley). *Description.* An essence-derived proportion is preserved at the cost of an RT60 that lands at 1.2 s in a deliberation room, because absorbing enough to reach 0.5 s would require covering a surface the designer wanted left hard for "proportional integrity". *Signature.* A measured RT60 above 0.8 s in a room intended for group discussion, with a design rationale referring to proportion. *Fix.* SA-10. **RT60 wins. Every time.** Chapter 14 established that the loudest speaker wins a reverberant room by physics; no proportional system compensates for that.

FM-31.7 THE BAND STRADDLING SCHROEDER

Damages: the coherence of the whole specification. *Description.* A folded band is chosen that spans `f_s`, so half the targets are modal and half are statistical, and the treatment strategy for the two halves is incompatible. *Signature.* A folded set spanning, say, 100-300 Hz in a room whose `f_s` is 147 Hz. *Fix.* SA-8. Compute `f_s` first, then choose the band, then fold.

FM-31.8 THE UNDISCLOSED LOUDSPEAKER

Damages: consent, indirectly but seriously. *Description.* The tuning is delivered by a ceiling speaker nobody knows about, at a level low enough that most occupants do not consciously notice it. Nothing is technically subliminal — it is above threshold — but nothing is disclosed either. *Signature.* An occupant who cannot say what is sounding in the room they work in. *Fix.* SA-2 and SA-13. **The disclosure requirement is not satisfied by audibility alone. A person must be able to know, not merely able in principle to notice.**

FM-31.9 THE PHANTOM SIX

Damages: the whole composition. *Description.* A twelve-house scheme is realised as twelve tuned tubes. Six of them are octave duplicates. The room contains twelve objects and six sounds, and the extra six are presented as distinct voices. *Signature.* Two tubes whose lengths stand in a ratio of exactly 2 or 4. *Fix.* SA-7. Build six, label each with the houses it carries — as §31.10's worked example does.

FM-31.10 THE UNBEATEN BEAT

Damages: the occupants' experience, and the credibility of the tuning. *Description.* PC5 and PC7 are sounded together, beat at 7.73 Hz, and the room is described as "harmonious". Occupants describe it as "buzzing" and nobody connects the two reports. *Signature.* Any published set with a folded pair under 20 Hz apart and no beat rate reported. *Fix.* SA-14. Report the beat. Then either separate the voices in time, transpose one by an octave, or accept the flutter and say so in the disclosure.

Worked Example

The Twelve-House Deliberation Room, carried end to end

A research institute organised into the twelve houses of Paper B §2.1 commissions a deliberation room for its cross-house committee. This is the same institute as Chapter 30's worked example; the room is where that committee will sit.

STEP 1 — Declare the category set and the ordinal rule.

The twelve houses, in the order Paper B publishes them, which is the declared ordinal:

n	House	Domain
1	Mathematics & Pattern	Formal structure, number, proof
2	Physics & Energy	Matter, force, measurement
3	Biology & Body	Life, organism, genetics
4	Psychology & Mind	Cognition, emotion, development
5	Consciousness & Field	Awareness, observation, qualia
6	Communication & Code	Value, exchange, accounting
7	Technology & Craft	Tools, systems, engineering
8	Spirituality & Practice	Contemplation, tradition, experience

n	House	Domain
9	Art & Expression	Aesthetics, creation, performance
10	Ecology & Earth	Ecosystem, soil, climate
11	Language & Story	Grammar, narrative, translation
12	Ethics & Governance	Law, sovereignty, distribution

Ordinal rule, as declared: the index is the house's position in the Paper B §2.1 enumeration. **This is a weak ordinal and must be reported as one.** Paper B is explicit that twelve "is not asserted as correct" and that the partition "is one idealised coding, not a uniquely privileged carving of the world". The enumeration order is not derived from any property of the houses. It is publishable, reproducible and stable — anyone with Paper B gets the same integers — but it is not intrinsic. **SA-5 requires this to be stated, and it is stated here.**

STEP 2 — Permutation sensitivity (ES-4). Three plausible re-orderings:

```
   swap 3 <-> 12   : both odd part 3  -> pitch-class displacement   0.00 cents
      swap  5  <->  7      :  class  5  <->  7      ->  displacement   772.63  -
737.65 =  34.98 cents
      swap  9  <->  11     :  class  9  <->  11     ->  displacement  1102.64  -
407.82 = 694.82 cents
```

Maximum displacement 694.82 cents — well over the 100-cent threshold. The ordering is load-bearing. Since the ordinal rule is weak (Step 1), this is the honest finding and it must travel with every number below: *a different but equally defensible enumeration of the twelve houses would produce a materially different tuning.* The room is being built to *an* essence set, not to *the* essence set.

STEP 3 — Raw essence frequencies. `f(n) = 0.483425732988414 * n^2`.

STEP 4 — Pitch classes by odd part. `N = 12`, so `ceil(12/2) = 6` classes.

STEP 5 — Fold into `[220, 440)` Hz. Band and `j` published per value (SA-6).

n	House	f(n) (Hz)	j	folded (Hz)	nearest note (A4=440)	dev (cents)	odd part	class
1	Mathematics & Pattern	0.483426	+9	247.5140	B3	+4.0	1	PC1
2	Physics & Energy	1.933703	+7	247.5140	B3	+4.0	1	PC1
3	Biology & Body	4.350832	+6	278.4532	C#4	+7.9	3	PC3
4	Psychology & Mind	7.734812	+5	247.5140	B3	+4.0	1	PC1
5	Consciousness & Field	12.085643	+5	386.7406	G4	-23.4	5	PC5

n	House	f(n) (Hz)	j	folded (Hz)	nearest note (A4=440)	dev (cents)	odd part	class
6	Communication & Code	17.403326	+4	278.4532	C#4	+7.9	3	PC3
7	Technology & Craft	23.687861	+4	379.0058	F#4	+41.7	7	PC7
8	Spirituality & Practice	30.939247	+3	247.5140	B3	+4.0	1	PC1
9	Art & Expression	39.157484	+3	313.2599	D#4	+11.8	9	PC9
10	Ecology & Earth	48.342573	+3	386.7406	G4	-23.4	5	PC5
11	Language & Story	58.494514	+2	233.9781	A#3	+6.6	11	PC11
12	Ethics & Governance	69.613306	+2	278.4532	C#4	+7.9	3	PC3

The six pitch classes, which are the entire harmonic vocabulary of this room:

class	houses carrying it	f(m) = k m^2 (Hz)	j	folded (Hz)	note	dev
PC1	Mathematics & Pattern; Physics & Energy; Psychology & Mind; Spirituality & Practice	0.483426	+9	247.5140	B3	+4.0
PC3	Biology & Body; Communication & Code; Ethics & Governance	4.350832	+6	278.4532	C#4	+7.9
PC5	Consciousness & Field; Ecology & Earth	12.085643	+5	386.7406	G4	-23.4

class	houses carrying it	f(m) = k m^2 (Hz)	j	folded (Hz)	note	dev
PC7	Technology & Craft	23.687861	+4	379.0058	F#4	+41.7
PC9	Art & Expression	39.157484	+3	313.2599	D#4	+11.8
PC11	Language & Story	58.494514	+2	233.9781	A#3	+6.6

Twelve houses, six voices — and the distribution is lopsided. Four houses share PC1, three share PC3, two share PC5, and PC7, PC9 and PC11 are singletons. **This is a real structural statement about the enumeration used: the houses at powers-of-two positions (1, 2, 4, 8) are harmonically indistinguishable from one another, and the houses at 7, 9 and 11 each have a voice of their own.** It is not a statement about mathematics, physics, psychology or spirituality; it is a statement about positions 1, 2, 4 and 8 in a list.

Sorted by pitch, with the intervals between adjacent classes:

#	class	folded Hz	note	step to next (cents)
1	PC11	233.9781	A#3 +7	97.4
2	PC1	247.5140	B3 +4	203.9
3	PC3	278.4532	C#4 +8	203.9
4	PC9	313.2599	D#4 +12	329.8
5	PC7	379.0058	F#4 +42	35.0
6	PC5	386.7406	G4 -23	330.0 (wrapping to 2x PC11)

Note that **PC11 has folded to a position *below* PC1**, because `121/128 = 0.945` is less than one. Folding into a fixed band reorders. The class labels are still correct; the ordering is not the index ordering, and comparing folded values across systems is FM-30.4. Within one room, though, this is simply the scale.

STEP 6a — Choose the room proportion (§31.6). Exhaustive search over class triples and octave offsets against the published ratio heuristics gave PC7 : PC9 : PC11 against Volkmann as the best match, error 0.031. Anchoring the axial fundamentals at the raw essence frequencies:

```
height  Lx = c / (2 * f(11)) = 343 / (2 * 58.494514) = 2.9319 m
width   Ly = c / (2 * f(9))  = 343 / (2 * 39.157484) = 4.3798 m
length  Lz = c / (2 * f(7))  = 343 / (2 * 23.687861) = 7.2400 m

ratio 1 : 1.4938 : 2.4694,  exactly 1 : 121/81 : 121/49
Volkmann heuristic 1 : 1.5 : 2.5     deviation −0.0062, −0.0306

V = 92.969 m^3
floor area = 31.709 m^2
total surface S = 131.555 m^2
occupancy at 2.5 m^2/person = 12.7 people
```

Twelve houses; twelve-and-a-bit seats. The room is the size of its own committee, and that was not arranged — it fell out of `c/(2f)`.

Cross-reference to Chapter 17: a room 4.38 m wide and 7.24 m long is a poor fit for a round table (the aspect ratio is 1.65) and a good fit for a long oval or an elongated polygon. **Chapter 17's requirement is sightline reciprocity — everyone sees everyone — and at 7.24 m the far-end sightline is at the outer limit of Hall's social distance (Chapter 12) and past comfortable facial-expression reading.** The honest resolution is to seat the twelve within the middle 5.5 m of the length and use the end 1.7 m for circulation and the monochord (Step 6d). **The proportional system proposed the room; Chapter 17's governance geometry disposes of how it is furnished.**

STEP 6b — Acoustic specification (§31.5, SA-10).

```
    Chapter 14 deliberation-room target: RT60 = 0.5 to 0.6 s

     RT60 = 0.5 s   ->   A = 0.161 * 92.969 / 0.5 = 29.94 m^2 sabins   -
>  alpha_bar = 0.2276
     RT60 = 0.6 s   ->   A = 0.161 * 92.969 / 0.6 = 24.95 m^2 sabins   -
>  alpha_bar = 0.1896

   Schroeder frequency at RT60 = 0.5 s:  f_s = 2000 * sqrt(0.5/92.969) = 146.7 Hz
```

Check SA-8: the folded set spans 233.98 to 386.74 Hz, entirely above `f_s = 146.7 Hz`. The whole set is in the statistical region and Sabine governs for all of it. The band is coherent.

Band priority from §31.5: four of six classes in the 250 Hz octave band, two in the 500 Hz band. **Specification: the treatment budget is weighted toward the 250 Hz band, where porous absorbers are losing efficiency — meaning at least one tuned membrane or Helmholtz element, not carpet alone.** Step 6d supplies exactly that.

STEP 6c — Cadence (§31.7, SA-11). The committee is a cross-house body; its cadence is taken from PC1, the class carried by the largest number of houses.

```
    f(PC1) = 0.483426 Hz  ->  T = 2.068570 s
    folded into [7, 14) days: p = 19, T * 2^19 = 12.5524 days

    Institute delta, measured (Chapter 24): 0.05 per week
    tau = 1/delta = 20 weeks
    P = 12.5524 days = 1.7932 weeks
    P/tau = 1.7932 / 20 = 0.0897
```

`P/tau = 0.090`, far below the 0.5 requirement of SA-11. Admissible without folding down. The forcing period is well inside the filter's time constant, so `Q` integrates smoothly rather than sawtoothing between meetings.

Second cadence, for the sub-committee on Technology & Craft (PC7):

```
    f(PC7) = 23.687861 Hz -> T = 0.042216 s -> p = 24 -> 8.1975 days
    P/tau = 8.1975/140 = 0.0586        (140 days = 20 weeks)
```

Beat between the two cadences: `12.5524 * 8.1975 / (12.5524 - 8.1975) = 23.63` days. The two bodies coincide about every 24 days — roughly monthly — and that coincidence

is a scheduling fact produced by the periods, with no coordinator required (Chapter 21).

Rounding, disclosed (SA-12): in practice both cadences will be rounded to the working calendar. 12.5524 days becomes a fixed twice-monthly slot (average 15.2 days, a departure of `1200*log2(15.2/12.5524) = 331` cents) and 8.1975 days becomes a fixed weekly-plus-one slot (8 days, a departure of 42 cents). **Those are the actual cadences. The derived values are the design intent and the rounding is declared. Anything else would be FM-31.5.**

STEP 6d — Tuned physical elements (§31.8, SA-13).

One six-string monochord, mounted along the north wall, steel strings `mu = 0.0012 kg/m` at 400 N (`v = 577.35 m/s`):

```
PC11  233.9781 Hz  ->  L = 1.2338 m
PC1   247.5140 Hz  ->  L = 1.1663 m
PC3   278.4532 Hz  ->  L = 1.0367 m
PC9   313.2599 Hz  ->  L = 0.9215 m
PC7   379.0058 Hz  ->  L = 0.7617 m
PC5   386.7406 Hz  ->  L = 0.7464 m
```

Overall instrument length about 1.40 m. **Each string carries an engraved label: its class, the houses it represents, its frequency to four figures, its essence index, and the sentence required by SA-15.**

One Helmholtz absorber tuned to PC11, neck radius 25 mm, neck length 100 mm, cavity 12.00 litres, resonating at 58.4945 Hz. **This sits within 0.005 Hz of the room's `(1,0,0)` axial mode at 58.494514 Hz (Step 6a), so it is a working modal trap for the height axis as well as a tuned element. It is the cheapest single item in the specification and the most acoustically useful.**

Beat report (SA-14): PC7 and PC5 differ by 7.7348 Hz and PC11 and PC1 by 13.5359 Hz. **Both are in the roughness zone. If all six strings are struck together the result flutters at about 8 Hz and again at about 13 Hz.** This is disclosed on the instrument's label. Occupants who want a clean chord strike PC1, PC3 and PC9 — the `1 : 9/8 : 81/64` Pythagorean triad, all within 8 cents of twelve-tone equal temperament.

STEP 7 — The disclosure notice (SA-1, SA-2, SA-3, SA-15). Posted at the entrance, reproduced here in full because it is part of the specification and not an appendix to it:

```
THIS ROOM IS PROPORTIONALLY TUNED. NOTHING IN IT IS HIDDEN.

Dimensions:  2.9319 m high, 4.3798 m wide, 7.2400 m long.
These were chosen so the room's three axial acoustic modes fall at
58.4945 Hz, 39.1575 Hz and 23.6879 Hz. The ratio is exactly 1 : 121/81 : 121/49.

A six-string monochord on the north wall is tuned to
233.98, 247.51, 278.45, 313.26, 379.01 and 386.74 Hz.
It sounds only when someone plays it. Anyone may play it. Anyone may stop.
Struck together, strings 5 and 6 beat at about 8 Hz and strings 1 and 2 at
about 13 Hz. This is expected.

A 12-litre Helmholtz absorber behind the north panel is tuned to 58.49 Hz.
It is passive. It makes no sound of its own.
```

```
 There is no loudspeaker in this room and no continuous tone of any kind.
There is no masked, hidden, subliminal, or below-threshold content of any kind
 in this room. If you can't hear it, it isn't here.

Derivation: the twelve houses of the institute, indexed by their position in the
 Paper B section 2.1 enumeration, through f(n) = k n^2 with
 k = 0.483425732988414 Hz. Twelve houses yield six distinct pitch classes.
 The full derivation is in the room file and anyone may have a copy.

These are ESSENCE frequencies, not physical resonances. No claim is made that
any material, room, institution or person physically resonates at these values.
The room's acoustic modes are ordinary acoustics: a room 4.3798 m wide has a
 mode at 39.16 Hz because sound travels at 343 m/s, and for no other reason.

The indexing of the twelve houses is not intrinsic. A different but equally
defensible ordering would give a different tuning. This is one essence set,
 not the essence set.
```

STEP 8 — What would falsify a claim about this room (SA-16, Chapter 25). No claim is made that the tuning improves anything. If the institute wishes to make one, the protocol is fixed in advance: measure turn-taking equality (Woolley's driver, Chapter 18) in the old room and the new one, with the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) − C_t + eps` fitted on the first eight committee cycles and frozen, reported on cycles 9-12 against all four required baselines — persistence, direct-input, matched autoregressive, and **the shuffled-input control**.

The prediction that would matter, and the one that is hardest to get: turn-taking equality improves in the new room *and* the improvement survives the shuffled-input control. If it does not survive, the honest conclusion is that the room is a good room because its RT60 is 0.5 s and its sightlines work, and that the essence tuning contributed a defensible way of choosing among many acceptable proportions and nothing further.

That would be a perfectly good result, and this manual would report it as one. Chapter 29's opening constraint holds all the way to the end: **the essence law generates targets. It does not generate effects, and this chapter has not claimed that it does.**

Chapter 32 . The Conversion Keys

Volume One engineered the outside — rooms, tables, rosters, institutions — with a single accumulator equation and four parameters. Volume Two turns the same instrument inward, and the first thing an inward-turned instrument needs is a set of units.

This chapter supplies them. It is the foundation chapter of Volume Two and everything after it depends on the arithmetic established here. **Four quantities this framework treats as separate subjects — interaction surplus, musical pitch, sensory magnitude, and the lived moment — are all logarithmic, and therefore any one converts into any other by an exact constant.** There are five constants, derived below and verified numerically, and they are small enough to memorise.

The claim is narrow and it is worth stating its narrowness first. **The four registers are interconvertible as measures. They are not identical as phenomena.** A surplus of 1.37 nats can be written as 3.96 octaves without loss, exactly as a kilogram can be written as 2.2046 pounds. Nothing about the pound is claimed to be a property of the kilogram. The keys are a change of units and nothing more — and a change of units is exactly what has been missing.

The Law

32.1 The four registers

Four quantities in this corpus are logarithmic. Their statuses differ and each is badged.

```
1.  SURPLUS      F     = ln N_eff = ln(1 + (N−1)u)      Paper A, Thm 2.1
                                                  [PROVEN, given axioms S1–S4]

2.  PITCH     p    = log2(f / f_ref)           [DEFINITION, universal in music]

3.  SENSATION    S     = k_F ∗ ln(I / I_0)              Fechner 1860
                                              [MEASURED LAW, with known limits]

4.  MOMENT      ln M  = ln D − ln P + ln E           log form of M = (D/P).E
                                             [POSTULATE − author's formula]
```

Register 1 is a theorem given four axioms, one of which (S3, affine effective count) was chosen and not derived. Register 2 is a definition. Register 3 is a nineteenth-century experimental result replicated continuously for a hundred and sixty years, with known failure regions at the extremes of intensity. Register 4 is **the author's postulate**, badged here and everywhere it appears, and subjected to the full Formula Validation Protocol in §32.13 and again in Chapter 34.

Mixing statuses in one table is what lets a postulate borrow the authority of a theorem. It will not be permitted here. **Every conversion in this chapter carries the status of the quantity it converts, and a conversion never upgrades a status.** A postulate converted into octaves is a postulate in octaves.

32.2 The only identity required

All four registers are logarithms. Logarithms in different bases differ by a constant factor and by nothing else:

```
log_b(x) = ln(x) / ln(b)

therefore, for any two bases b1, b2:

log_{b2}(x) = log_{b1}(x) * ( ln(b1) / ln(b2) )
```

That is the whole mechanism. **Because every register is a logarithm, every conversion between registers is a multiplication by a constant, and no information is lost, gained, or distorted in the crossing.** There is no approximation in this chapter's unit arithmetic; the only approximations are in the measured inputs. Verified numerically at 16 significant figures:

```
 ln(3)   = 1.0986122886681098
log2(3) = 1.5849625007211563      ln(3)/ln(2) = 1.5849625007211563   agree
 ln(144) = 4.969813299576001
log2(144) = 7.169925001442312    ln(144)/ln(2) = 7.169925001442313  agree to 1 ulp
```

32.3 The five constants

```
ln 2      = 0.6931471805599453        nats per bit
1 / ln 2  = 1.4426950408889634        bits per nat
2 / ln 2  = 2.8853900817779268        ESSENCE OCTAVES per nat
2400/ln 2 = 3462.4680981335123        cents per nat
k         = 0.4834257329884143 Hz     the ground essence frequency
```

The fifth is not a unit conversion; it is the reference pitch, and it is **a chosen constant, not a measured one** (CANON_SONIC §1). Its value is `k = (9/8)^2 / phi^2 = (81/64)/(phi+1)`, a Pythagorean ditone divided by the golden section. Verified:

```
phi   = 1.618033988749895
phi^2 = 2.618033988749895        = phi + 1, as required
(9/8)^2 / phi^2   = 0.48342573298841435
(81/64) / (phi+1) = 0.48342573298841435    identical
```

32.4 The conversion key table

Read a row as "one of this unit equals", and a column as "measured in this unit".

1 of v	nats	bits	octaves	cents
nat	1	1.442695	2.885390	3462.468098
bit	0.693147	1	**2 exactly**	**2400 exactly**
octave	0.346574	**0.5 exactly**	1	**1200 exactly**
cent	0.000288811	0.000416667	0.000833333	1

Three entries are exact, and the reason is not decorative. The 1200 is definitional: the cent has

been one twelve-hundredth of an octave since Ellis. The 2 and the 0.5 come from somewhere else entirely, and §32.6 derives them.

32.5 The derived conversion laws

```
SURPLUS -> PITCH      octaves above ground = 2F / ln 2 = 2.885390 * F     (F in nats)

 SURPLUS -> FREQUENCY     f_cross = k * e^(2F) = k * N_eff^2

ESSENCE -> FREQUENCY    f(n)    = k * n^2                 n = essence index

 INTENSITY -> SENSATION  S       = k_F * ln(I / I_0)         Fechner

INTENSITY -> MAGNITUDE  log psi = log k_S + a * log I     Stevens; a IS the key

MOMENT -> SURPLUS     M normalised onto [0, ln N] serves as S_t in the storage law
                           [NORMALISATION CHOSEN, not derived]

DISTORTION          Delta = |ln r|,  r = D/P        reported in any register
```

32.6 Why one bit is two octaves and not one

This is the single non-obvious constant in the table and it must be derived rather than asserted, because a reader who assumes the ordinary information-theoretic correspondence will get it wrong by a factor of two.

One bit of surplus means the effective count doubled: `N_eff -> 2 N_eff`. The essence law (CANON_SONIC §1) maps index to frequency by a **square**:

```
f(n) = k n^2

f(2n) / f(n) = (2n)^2 / n^2 = 4      =  two octaves, exactly
```

Doubling the effective count quadruples the essence frequency, and quadrupling is exactly two octaves. That, and only that, is why one bit is two octaves in this book. The factor two in `octaves = 2F / ln 2` is the exponent of the essence law and nothing else. Verified across the range:

```
N_eff  ->  2*N_eff        f ratio      octaves
  1    ->    2             4.0          2.0
  3    ->    6             4.0          2.0
 12    ->   24             4.0          2.0
```

The consequence for the honest reader is a caution. If the essence law's exponent were some other value `s`, so that `f = k n^s`, the key would be `s / ln 2` octaves per nat:

exponent s	octaves per nat	cents per nat
0.5	0.721348	865.617
1	1.442695	1731.234
1.5	2.164043	2596.851

exponent s	octaves per nat	cents per nat
2 (this book)	**2.885390**	**3462.468**
2.5	3.606738	4328.085
3	4.328085	5193.702

The surplus-to-pitch key is therefore contingent on the square law, and the square law is a postulate about essence structure, not a measured fact (CANON_SONIC §6). State this whenever 2.885390 is used. The constant is exact given the postulate and not exact independently of it — the difference between "2.885390 is the correct number" and "2.885390 is a correct number about the world." Only the first is claimed.

32.7 The honesty clauses, restated in full

Binding wherever this chapter's material is used, and reproduced rather than cross-referenced because this is the chapter in which the temptation to over-read is maximal.

- **The output of the essence law is an ESSENCE FREQUENCY, not a physical resonance.** No claim is made that any material, molecule, element, room, institution or person physically resonates at these values. A surplus of 1.37 nats converts to 120.07 Hz in the sense that 20 degrees Celsius converts to 68 degrees Fahrenheit. Nothing is vibrating.
- **k is a chosen reference constant**, a design decision by the author. Badge it every time.
- **The square law f = k n^2 is a postulate about essence structure.** Its formal coincidence with the eigenvalue spacing of a one-dimensional infinite square well, where E_n is proportional to n^2, is a genuine mathematical correspondence and may be noted as such — and coincidence of form is not evidence of physical mechanism, which must be said in the same breath.
- **The essence index requires a defensible ordering.** An arbitrary ordinal yields arbitrary frequencies. Chapter 33 is entirely about this requirement.
- **Surplus, as deployed, is a LOWER BOUND on the ensemble's information.** We log the floor, not the ceiling. A lower bound converted into octaves is a lower bound in octaves.
- **Paper A is pure mathematics and makes no empirical or metaphysical claim.** Paper D does not identify surplus with Shannon entropy or mutual information in general. Paper E proves no thermodynamic negentropy law.
- **Stevens' exponents are constants of a procedure, not of nature.** They vary with method, range and instruction.
- **The Moment equation is the author's postulate.** Every conversion involving it inherits that status.

The Engineering

32.8 Deriving each constant

Nats to bits. A nat is the logarithm base e, a bit base 2. By §32.2, one nat is `1/ln 2` bits.

```
1 / ln 2 = 1 / 0.6931471805599453 = 1.4426950408889634
```

To six places, 1.442695. Bits are the register in which an engineering reader has the strongest intuitions: `ln 12 = 2.484907` nats is `3.584963` bits, so a twelve-category system, fully orthogonal, is worth three and a half bits of surplus per interaction. That is a small number and it should feel small.

Nats to octaves. By §32.6 one unit of F multiplies frequency by e^2, and the octaves in a ratio R are `log2(R)`. So

```
octaves per nat = log2(e^2) = 2 / ln 2 = 2.8853900817779268
```

To six places, 2.885390, exactly as CANON_INTERNAL §1 records.

Nats to cents. A cent is one twelve-hundredth of an octave by definition, so

```
cents per nat = 1200 * (2 / ln 2) = 2400 / ln 2 = 3462.4680981335123
```

To four places, 3462.4681. CANON_INTERNAL §1 records 3462.4681. Verified.

Octaves to nats. The reciprocal: `ln 2 / 2 = 0.3465735902799726`4, 0.346574 to six places.

Cents to nats. `ln 2 / 2400 = 0.00028881132523331055` — the resolution constant. One cent corresponds to about 289 millionths of a nat of surplus. **The pitch register has far finer graduation than the surplus register has meaningful precision**, which is a warning and not an invitation: see F-32.7, the Cents Illusion.

32.9 The surplus-to-pitch conversion in use

The working table of Volume Two converts an effective count into all four registers plus the essence frequency. It is computed, not transcribed.

N_eff	F (nats)	bits	octaves	cents	f = k.n² (Hz)
1	0.000000	0.000000	0.000000	0.000	0.483426
2	0.693147	1.000000	2.000000	2400.000	1.933703
3	1.098612	1.584963	3.169925	3803.910	4.350832
4	1.386294	2.000000	4.000000	4800.000	7.734812
5	1.609438	2.321928	4.643856	5572.627	12.085643
6	1.791759	2.584963	5.169925	6203.910	17.403326
7	1.945910	2.807355	5.614710	6737.652	23.687861
8	2.079442	3.000000	6.000000	7200.000	30.939247
9	2.197225	3.169925	6.339850	7607.820	39.157484
10	2.302585	3.321928	6.643856	7972.627	48.342573
12	2.484907	3.584963	7.169925	8603.910	69.613306
16	2.772589	4.000000	8.000000	9600.000	123.756988
20	2.995732	4.321928	8.643856	10372.627	193.370293
24	3.178054	4.584963	9.169925	11003.910	278.453222

Two features are worth dwelling on. **The powers of two are exact in three registers at once:** `N_eff = 2, 4, 8,` 16 give exactly 1, 2, 3, 4 bits; exactly 2, 4, 6, 8 octaves; exactly 2400, 4800, 7200, 9600 cents. Nothing is rounded.

And the frequency column is brutal. Going from an effective count of 2 to 24 — a twelvefold increase in genuinely different views — moves the surplus from 0.693 to 3.178 nats, a factor of 4.6, but moves the essence frequency from 1.93 Hz to 278.45 Hz, a factor of 144. **The pitch register amplifies exactly the differences the surplus register compresses, which is why it is useful as a display register and dangerous as an inference register.** Use frequency to hear a system, not to argue about it.

A fifth conversion key, and it converts surplus into time. The four registers of §32.1 are all instantaneous. Chapter 71 adds one that is not, derived from the Gabor time-bandwidth limit `dt * df >= 1/(4 pi)` together with this chapter's own relation `d(ln f) = 2 dF`:

```
SURPLUS DIFFERENCE -> TIME    cycles required >= 1 / (8 pi * dF)    1/(8 pi) = 0.039789
```

Resolving a difference dF in surplus costs at least `1/(8 pi dF)` cycles of the rendered tone, and the constant is exact. Two properties make it a key in the sense of §32.1 rather than a rule of thumb. It is **independent of the transposition**, because it constrains a ratio and octave transposition multiplies numerator and denominator alike — the same cycle count serves at 30 Hz and at 19 kHz. And it is **denominated in cycles, not seconds**; converting to wall-clock requires the frequency, which is exactly the dependence you want kept visible.

dF (nats)	cycles required	at `N_eff = 12` (69.613 Hz)	at `N_eff = 24` (278.453 Hz)
0.10	0.398	5.72 ms	1.43 ms
0.05	0.796	11.43 ms	2.86 ms
0.01	3.979	57.16 ms	14.29 ms
0.005	7.958	114.31 ms	28.58 ms
0.001	39.789	571.57 ms	142.89 ms

Two boundary conditions belong with the table. **The frequency column has a working range and it is short.** Under a single integer octave transposition into 20 Hz - 20 kHz the ladder fits only at `j = 6`, and only up to **`N_eff = 25`**: `k * 25^2 * 64 = 19337.0` Hz, while `N_eff = 26` lands at 20914.9 Hz and leaves the band. **Above an effective count of 25 the pitch register cannot be rendered in one piece and requires split transposition** — a declared design choice, not a formatting detail. And **the rectangular gate is disqualified as an envelope**: it has no finite spectral second moment, so a tone rendered with a hard on/off gate has no defined bandwidth and the key above does not apply to it. Window the envelope; Chapter 31 §31.9 gives the default.

32.10 The consistency check that ties the two frameworks together

CANON_SONIC §4 asserts that the effective count is the cross-disciplinary essence index. That assertion has an arithmetical consequence which can be checked and which, if it failed, would break the bridge:

```
within a discipline:      f = k * n^2       n an integer index
across disciplines:       f = k * e^(2F)    F = ln N_eff

at u = 1 (fully orthogonal), N_eff = N, so F = ln N, so

k * e^(2 ln N) = k * N^2                    the two must coincide
```

Verified to machine precision:

n	k.e^(2.ln n)	k.n²	difference
1	0.48342573298841435	0.48342573298841435	0
2	1.9337029319536574	1.9337029319536574	0
3	4.35083159689573	4.350831596895729	8.9e-16
7	23.6878609164323	23.687860916432303	3.6e-15
12	69.61330555033166	69.61330555033166	0
24	278.4532222013267	278.45322220132664	5.7e-14

Residuals are floating-point rounding at the sixteenth significant figure, not model error.

The two frameworks agree exactly at the orthogonal limit, which means the essence index of a fully non-parallel N-category system is N itself. A discipline with twelve mutually orthogonal categories sits at index 12, essence frequency 69.613306 Hz, 7.169925 octaves above ground — precisely the figure CANON_SONIC §4 records. Twelve non-orthogonal categories sit at a non-integer index between 1 and 12, and that is the crucial difference between the two uses of the law:

Within a discipline the index is an integer and the lattice is discrete. Across disciplines the index is the effective count, which is continuous, and the pitch is continuous with it. The discrete lattice is Chapter 33. The continuous bridge is here. They meet only at `u = 1`.

Worked, for a twelve-category system at various levels of non-parallelism:

u	N_eff	F (nats)	octaves	f (Hz)
0.00	1.000	0.00000	0.0000	0.4834
0.25	3.750	1.32176	3.8138	6.7982
0.50	6.500	1.87180	5.4009	20.4247
0.75	9.250	2.22462	6.4189	41.3631
1.00	12.000	2.48491	7.1699	69.6133

Read the first row. **Twelve parallel categories are the ground essence: 0.483426 Hz, zero octaves, zero surplus.** The pseudo-diversity failure of Chapter 11 has a pitch, and its pitch is the bottom of the instrument.

FVP . The Register Conversion Identity

```
x_target = x_source * ( ln b_source / ln b_target )
```

Gate	Test	Result
FVP-1 DIMENSIONS	Both sides are dimensionless logarithms of a ratio; the multiplier is a ratio of logarithms of pure numbers. Units agree.	**PASS**
FVP-2 SCALE TYPE	All four registers are ratio-scaled *in the argument* and interval-scaled *in the logarithm*: a difference of two pitch values is meaningful, an absolute value requires a stated reference. The conversion is a positive linear map, permissible on interval and ratio scales alike.	**PASS, reference required in print**
FVP-3 LIMITING CASES	`x -> 0` gives 0 in every register (unison, zero surplus, ground essence). `x -> infinity` diverges in every register at the same rate. No register saturates or wraps.	**PASS**
FVP-4 FIXED POINTS	`x = 0` is the unique fixed point: zero surplus is zero bits is zero octaves is zero cents. **The ground is shared by all four registers**, which is what makes them comparable.	**PASS**
FVP-5 IDENTIFIABILITY	The multiplier is fixed by the two bases; no free parameters to recover.	**N/A — no free parameters**
FVP-6 FALSIFICATION	Not empirically falsifiable: an identity of logarithms, arithmetically checkable and checked (§32.2, §32.10). The framework claim riding on it — that the essence exponent is 2 — **is** falsifiable, and is flagged in §32.6.	**N/A (identity); dependent postulate flagged**
FVP-7 REPRODUCIBILITY	Any correct implementation returns the same constants. Verified in double precision and cross-checked in 40-digit decimal arithmetic; `ln 2` agreed to all 40 digits.	**PASS**

32.11 The sensation register: Weber, Fechner, and the just-noticeable difference

The third register settles the question a sceptical reader brings to Volume Two: **can a subjective quantity be measured at all?** Yes, and it has been done since 1834. Volume Two's entire warrant rests on the method, so the method is stated precisely.

Weber (1834): the just-noticeable difference in a stimulus is a constant *ratio* of the stimulus, not a constant amount.

```
dI / I = c        c = the Weber fraction for that modality and range
```

Fechner (1860): integrate Weber. If each just-noticeable difference is one subjective unit and each costs a fixed proportional increase, then counting units from a reference gives

```
S = k_F * ln( I / I_0 )
```

Sensation is logarithmic in intensity. **This is why the framework is logarithmic: not because logarithms are elegant, but because the register perception itself operates in is logarithmic and has been known to be for a hundred and sixty years.**

The number of just-noticeable differences between a reference and a stimulus is

```
n_JND = ln( I / I_0 ) / ln( 1 + c )
```

A JND is itself a logarithmic unit, so it joins the key table. The following is **parametric, not measured**: it asserts no value of c for any modality, because Weber fractions are modality-, range- and method-dependent and this manual does not report figures it has not measured or cited.

c	nats/JND	bits/JND	octaves/JND	cents/JND	JNDs across x10
0.01	0.009950	0.014355	0.028711	34.453	231.4
0.02	0.019803	0.028569	0.057138	68.566	116.3
0.05	0.048790	0.070389	0.140779	168.934	47.2
0.10	0.095310	0.137504	0.275007	330.008	24.2
0.15	0.139762	0.201634	0.403268	483.921	16.5
0.20	0.182322	0.263034	0.526069	631.283	12.6
0.30	0.262364	0.378512	0.757023	908.428	8.8

A modality with a Weber fraction of 0.10 has a JND worth 0.0953 nats. Ten JNDs is 0.953 nats, less than the surplus ceiling of a three-category system. That is the kind of statement the conversion keys make available and that is impossible without them.

Honest limits, stated with the law: Fechner's law fails near absolute threshold and near the upper end of a modality's range, where the Weber fraction is not constant. The logarithmic form is an excellent middle-range description, not a universal law of sensation. Stevens' power law is the competitor and in many modalities the better fit.

FVP . Fechner's Law

```
S = k_F . ln( I / I_0 )
```

Gate	Test	Result
FVP-1 DIMENSIONS	`I/I_0` is dimensionless by construction, so the logarithm is admissible. `S` carries the units of `k_F`, the scale constant of the chosen sensation unit (JNDs, or nats via §32.8).	**PASS if I_0 declared; FAIL if the reference is omitted**
FVP-2 SCALE TYPE	`I` is ratio-scaled with a true zero. `S`, built by JND-counting, is at best **interval-scaled**: JND counting gives equal-appearing increments, not verified equal ratios. This is the classic objection to Fechner and it is a real one. Ratio-scale sensation requires magnitude estimation, not JND integration.	**PARTIAL — S is interval; do not divide two Fechner sensations**
FVP-3 LIMITING CASES	`I -> I_0` gives `S -> 0`. `I -> 0` gives `S -> -infinity`: the divergence is the absolute threshold, below which the law is known to fail. `I -> infinity` gives unbounded `S`, failing at every modality's damage ceiling. **Both divergences name real phenomena and mark the law's known boundaries.**	**PASS, both boundaries named**
FVP-4 FIXED POINTS	`I = I_0` is the unique invariant: the scale has no content until a reference is chosen, and the reference is a decision.	**PASS**
FVP-5 IDENTIFIABILITY	`k_F` and `I_0` trade off — scaling `I_0` shifts `S` by an absorbable constant. **Identifiable only up to that offset**, so a fit must fix `I_0` by independent operational definition (measured absolute threshold).	**CONDITIONAL — fix I_0 externally or the fit is degenerate**

Gate	Test	Result
FVP-6 FALSIFICATION	Falsified in a modality if the measured Weber fraction varies systematically with `I` across the working range. It has been so falsified at the extremes: dead as a universal law, standing as a middle-range description.	**PASS — and has partially failed, honestly**
FVP-7 REPRODUCIBILITY	Weber fractions recover across subjects and laboratories within modality-specific bounds — among the most reproducible findings in psychology. **This gate is why psychophysics is the existence proof for Volume Two.**	**PASS**

32.12 Stevens' exponents are measured conversion keys

Stevens (1957) proposed the power law and, more importantly, the method:

```
psi = k_S * I^a
```

Take logarithms of both sides:

```
log psi = log k_S + a * log I
```

In log-log coordinates the power law is a straight line, and its slope is a. The exponent is therefore a conversion key in exactly the sense of this chapter: it is the number of perceptual log-units bought by one physical log-unit. It is not a fitted nuisance parameter. It is the constant of the change of units between the physical register and the perceptual one, and it is obtained by measurement rather than by definition, which is what distinguishes it from every other constant in §32.3.

There is a way of saying this that makes it immediately usable:

```
perceptual octaves per physical octave = a       exactly
```

An octave is a doubling in either register, and doubling `I` multiplies `psi` by `2^a`, which is a octaves. The exponent *is* the octave key. Published, replicated exponents, nominal and method-dependent:

Modality	a	world x2 -> psi x	world x10 -> psi x	psi x2 needs world x	cents per physical octave
Loudness (1 kHz)	0.67	1.5911	4.6774	2.8138	804.0
Brightness	0.33	1.2570	2.1380	8.1698	396.0
Visual length	1.00	2.0000	10.0000	2.0000	1200.0
Taste (salt)	1.30	2.4623	19.9526	1.7044	1560.0
Heaviness (lifted)	1.45	2.7321	28.1838	1.6129	1740.0
Warmth on arm	1.60	3.0314	39.8107	1.5422	1920.0
Electric shock	3.50	11.3137	3162.2777	1.2190	4200.0

`a < 1` is compressive; `a = 1` is linear; `a > 1` is expansive. **Most senses compress. Pain expands.** The consequence recurs throughout Volume Two: to double perceived loudness you must roughly triple the physical intensity (2.81x), but to double perceived shock you need only a 22% increase. **Comfort is expensive to increase and cheap to lose, and the exponents say so in numbers.**

The most consequential row is brightness: doubling perceived brightness requires **8.17 times** the physical light. A lighting specification written in lumens and read by an occupant in perceived brightness is a specification in two registers, mismatched by a factor of eight. Chapter 12 specified rooms; this is the arithmetic beneath it.

Cross-modality matching. Because each exponent is a key to the same physical register, the key between two modalities is the ratio of their exponents: if a subject adjusts modality *i* to match modality *j*, the predicted log-log slope is `a_j / a_i`. This is the strongest test of the power law, because the matching exponent is predicted in advance from two independently measured exponents and then checked.

matched ->	Loudness	Brightness	Length	Salt	Heaviness	Warmth	Shock
Loudness	1.0000	2.0303	0.6700	0.5154	0.4621	0.4188	0.1914
Brightness	0.4925	1.0000	0.3300	0.2538	0.2276	0.2062	0.0943
Length	1.4925	3.0303	1.0000	0.7692	0.6897	0.6250	0.2857
Salt	1.9403	3.9394	1.3000	1.0000	0.8966	0.8125	0.3714
Heaviness	2.1642	4.3939	1.4500	1.1154	1.0000	0.9062	0.4143
Warmth	2.3881	4.8485	1.6000	1.2308	1.1034	1.0000	0.4571
Shock	5.2239	10.6061	3.5000	2.6923	2.4138	2.1875	1.0000

The table is computed from the nominal exponents above; it is a set of **predictions**, not observations, and its honest use is as a preregistration target. **A framework that predicts a matching slope of 2.0303 for loudness-to-brightness and then measures 2.03 has earned something.**

A framework that fits the slope after the fact has earned nothing. That is FVP-6, and it is why the table appears here rather than in a results section.

Honest limits, restated: Stevens' exponents vary with method (estimation versus production give systematically different slopes — the regression effect), with stimulus range, and with instruction. **Report them as measured constants of a procedure, not constants of nature.**

FVP . Stevens' Power Law

`psi = k_S . I^a`, equivalently `log psi = log k_S + a . log I`

Gate	Test	Result
FVP-1 DIMENSIONS	`I` must be a ratio to a stated reference before exponentiation; `I^a` of a dimensioned quantity is ill-formed. `k_S` carries the units of `psi`.	**PASS if reference declared; a common silent failure otherwise**
FVP-2 SCALE TYPE	`I` ratio-scaled with true zero. `psi` **ratio-scaled** — the whole point of magnitude estimation, which asks subjects for ratios directly and validates them by cross-modality matching. Multiplication and division of `psi` are permissible.	**PASS — the gate Fechner fails and Stevens passes**
FVP-3 LIMITING CASES	`I -> 0` gives `psi -> 0` for all `a > 0`: no stimulus, no sensation. `I -> infinity` diverges; real modalities terminate at damage or saturation, so an upper working range must be declared. At `a = 1` the law degenerates to proportionality (visual length), as observed.	**PASS, upper bound declared per modality**
FVP-4 FIXED POINTS	`I = I_ref` gives `psi = k_S`. In log-log coordinates the invariant is the intercept: the anchor stimulus. **The anchor is chosen by the experimenter, which is exactly why instruction affects the exponent.**	**PASS, dependency named**

Gate	Test	Result
FVP-5 IDENTIFIABILITY	a is the log-log slope, k_S the intercept; separately recoverable from two or more levels spanning a decade. Cross-modality matching independently recovers the a ratios, over-identifying the system.	**PASS — over-identified, which is strong**
FVP-6 FALSIFICATION	Killed in a modality if log-log plots are reliably curved over the declared range, or if the matching slope departs from the predicted a_j/a_i beyond measurement error. Both tests have run for decades.	**PASS — a live falsification test**
FVP-7 REPRODUCIBILITY	Exponents recover across laboratories **within a fixed procedure** and fail to recover across procedures. Honest verdict: constants of a specified method, not of a modality. **Report the procedure or the number is not a measurement.**	**CONDITIONAL PASS — procedure-bound**

32.13 The moment register, and its conversion

The Moment equation is the author's postulate:

```
M = ( D / P ) . E        D perspective, P perception, E environment, M moment
```

Taking logarithms puts it in the same register as everything else in this chapter:

```
ln M = ln D − ln P + ln E
```

In the log register the Moment equation is additive, and that is its most useful form. A multiplicative model with a division becomes a sum of three terms, each reportable, budgetable and convertible independently. Full treatment is Chapter 34; the conversion arithmetic belongs here.

The calibration ratio and its distortion. With `r = D/P`, distortion is `Delta = |ln r|`, a log-ratio and therefore expressible in any register in the key table:

r = D/P	state	Delta (nats)	bits	octaves	cents
0.25	severely under-framed	1.386294	2.00000	4.00000	4800.00
0.50	under-framed	0.693147	1.00000	2.00000	2400.00
0.80	mildly under-framed	0.223144	0.32193	0.64386	772.63
1.00	**calibrated — contact**	**0.000000**	0.00000	0.00000	0.00
1.25	mildly over-framed	0.223144	0.32193	0.64386	772.63
1.80	over-framed	0.587787	0.84800	1.69599	2035.19
2.00	over-framed	0.693147	1.00000	2.00000	2400.00
4.00	severely over-framed	1.386294	2.00000	4.00000	4800.00

Note the symmetry, which is the reason for the logarithm: `r = 0.5` and `r = 2.0` are equally distorted, 2400 cents in either direction. **A raw ratio makes over-framing look worse than under-framing because it is unbounded above and bounded below by zero; the log-ratio makes them commensurate, and they should be commensurate, because both are the same failure of match.**

The moment as storage input. CANON_INTERNAL §1 records the join: `M` normalised onto `[0, ln N]` serves as `S_t` in the storage law. **The normalisation is a chosen mapping, badged here as a design decision.** It buys Chapter 40's self-storage law, in which the moment becomes the input that builds character.

FVP . The Moment Equation in Log Form

`ln M = ln D − ln P + ln E` — **[POSTULATE — the author's formula, badged]**

Gate	Test	Result
FVP-1 DIMENSIONS	`D` and `P` **must share units** so `r = D/P` is dimensionless; `M` inherits `E`'s units. `ln D − ln P` is well-formed iff `D` and `P` share both units and reference.	**CONDITIONAL — shared units are a requirement, not a result**

Gate	Test	Result
FVP-2 SCALE TYPE	Division requires **ratio scale with a true zero** for `D` and `P`. **This is the crux of the entire volume.** With only ordinal data the formula is a heuristic ordering, not an equation, and must be demoted in print. Chapter 53 is devoted to this gate.	**UNRESOLVED PENDING MEASUREMENT**
FVP-3 LIMITING CASES	`P -> 0`: `M -> infinity` — frame with no intake: hallucination, dogma, apophenia. `D -> 0`: `M -> 0` — sensation is not experience. `E -> 0`: `M -> 0`. **All three divergences name real, observable phenomena.**	**PASS — every divergence is nameable**
FVP-4 FIXED POINTS	`D = P` gives `M = E`. **Frame matches intake and the moment equals the environment: veridical contact, stated as an equation.** The formula's most important content and the exemplar for the whole FVP.	**PASS — fixed point is interpretable**
FVP-5 IDENTIFIABILITY	**FAILS as written.** `(D, P, E)` and `(cD, cP, E)` predict identical `M` for any `c > 0`; only `r` and `rE` are identified. The formula is testable in `r` and `M`, not in `D` and `P` separately, unless an independent operationalisation pins one scale.	**FAIL — identified only up to a common scaling of D and P**
FVP-6 FALSIFICATION	Template (Paper E §5.3): preregister an operationalisation of `r` and `M`; fit; freeze; report held-out; beat persistence, direct-input, matched autoregressive and **shuffled-input** baselines. Killed if held-out `M` is not predicted better than the shuffled-input control.	**DEFINED, NOT YET RUN**

Gate	Test	Result
FVP-7 REPRODUCIBILITY	Same operationalisation, different people, different laboratories: do `r` and the scale of `M` recover? **Not yet attempted. This gate decides whether the formula is science, and it is the one most often skipped.**	**NOT YET ATTEMPTED**

Two gates fail or stall, and both are printed as failures. **The identifiability failure is fixable, and it tells the researcher exactly what to do next: pin the scale of `P` by an independent measurement.** That is the FVP working correctly.

32.14 The forced consequence: the maximum possible moment is ln N

This is the result the chapter exists to reach. It is **derived, not chosen**, and the derivation is short.

Begin with the internal surplus (CANON_INTERNAL §8): Paper A applied within a single observer. A perceiver holds `N` distinguishable categories. The framings brought to bear on a moment are states in that category space, with mutual non-parallelism `u`. The internal effective count is

```
N_eff = 1 + (N − 1) u              u in [0,1]
F_int = ln N_eff
```

Paper D, Corollary 2.2 gives the range: `0 <= F <= ln N`, with `F = ln N` attained if and only if the framings are mutually orthogonal (Paper A, Theorem 4.3). Now apply the chosen normalisation of §32.13: **the moment, expressed in the surplus register, is `F_int`.** Therefore

```
0 <= (moment, in the surplus register) <= ln N

MAXIMUM POSSIBLE MOMENT = ln N
```

The richest moment available to a perceiver is bounded by the logarithm of the number of kinds of thing that perceiver can distinguish. No arrangement of attention, no intensity of experience, no quality of environment raises it. It is a ceiling on the perceiver, not on the world.

This is Ashby's Law of Requisite Variety appearing in first-person experience. Ashby (1956): only variety absorbs variety, and a regulator's capacity is bounded by its own variety. Here the regulator is a perceiving mind, the disturbance is the world, and the bound is `ln N`. **It was not chosen. It follows from the range theorem plus the normalisation, and the normalisation is the only chosen step.**

Three consequences, developed later in the volume.

First, the ceiling is logarithmic, so it is nearly immovable by effort. Doubling `N` — doubling the genuinely distinct categories a person can bring to bear — adds `ln 2 = 0.693147` nats, or exactly one bit, or exactly two essence octaves, **no matter where you start.**

N	ln N (nats)	bits	octaves	cents
2	0.693147	1.00000	2.00000	2400.00
4	1.386294	2.00000	4.00000	4800.00
8	2.079442	3.00000	6.00000	7200.00
16	2.772589	4.00000	8.00000	9600.00
32	3.465736	5.00000	10.00000	12000.00
64	4.158883	6.00000	12.00000	14400.00
128	4.852030	7.00000	14.00000	16800.00
256	5.545177	8.00000	16.00000	19200.00
1000	6.907755	9.96578	19.93157	23917.88

Going from 2 distinguishable categories to 1000 — a five-hundred-fold increase in vocabulary — buys 6.21 nats, less than nine bits. This is the internal form of Chapter 8's result, with the same moral: you cannot escape a capacity problem by adding categories.

Second, `N` is trainable — slowly, logarithmically, and only by acquiring genuinely distinct categories rather than relabelled ones. That is what an education is, stated as an inequality.

Third, the ceiling is rarely the binding constraint. Most perceivers are nowhere near `ln N`, because u is low: the framings they bring are near-parallel. **The gain from raising u is almost always far larger than the gain from raising `N`.** Book One's Node 03 HORIZON, the many windows of the one tower, is this statement in narrative form; Chapter 39 derives it properly.

FVP . The Maximum Moment

`M_max = ln N` (moment expressed in the surplus register)

Gate	Test	Result
FVP-1 DIMENSIONS	Both sides in nats; `N` a dimensionless count. Consistent.	**PASS**
FVP-2 SCALE TYPE	`N` is a **count** — ratio-scaled with a true zero — provided "distinguishable category" has an operational definition. Without one, `N` is not measured and the bound is not usable. **The gate transfers to the operationalisation of category count: Chapter 33's subject.**	**CONDITIONAL — requires a defensible category set**

Gate	Test	Result
FVP-3 LIMITING CASES	`N = 1`: `ln 1 = 0`. A perceiver with one category has no moment available — correct and severe. `N -> infinity`: the bound diverges logarithmically, arbitrarily slowly. No finite perceiver approaches it.	**PASS**
FVP-4 FIXED POINTS	Attained iff `u = 1`, the framings mutually orthogonal (Paper A Thm 4.3). **The invariant is orthogonality and it is unreachable in practice** — the honest content of the bound: a ceiling, not a target.	**PASS**
FVP-5 IDENTIFIABILITY	`N` and `u` are separately recoverable if category membership and pairwise non-parallelism are separately measurable; confounded if `u` is estimated by a procedure that already assumes the category count. **Estimate `u` by error decorrelation, not by labels.**	**PASS with a named confound**
FVP-6 FALSIFICATION	The bound is a theorem given the axioms plus the normalisation, so not falsifiable *as arithmetic*. The **normalisation** is: if a magnitude-estimation moment scale exceeded `ln N` for a subject whose `N` was independently established, the normalisation is wrong. A runnable experiment.	**NORMALISATION FALSIFIABLE; TEST SPECIFIED**
FVP-7 REPRODUCIBILITY	Requires `N` to recover across instruments for the same person — the specific-objectivity requirement of Chapter 54, with Rasch separability as the standard. Not yet attempted.	**NOT YET ATTEMPTED**

32.15 What the keys do not license

A conversion table is an invitation to one class of error, and the class deserves a name. **The keys convert measures. They do not transport warrant.**

- A **lower bound** converted is a lower bound, not an estimate. "3.96 octaves" means "at least 3.96 octaves" and must be written that way (CANON §1.1).
- A **postulate** converted is a postulate. The Moment equation in cents is still the author's postulate.
- A **procedure-bound measurement** converted is procedure-bound. A loudness exponent of 0.67 written as 804 cents per physical octave is still 0.67 by that method, range and instruction.
- A surplus converted into a **frequency** is an essence frequency. **Nothing is vibrating.** This is the error the chapter most expects and the one the house rules most firmly forbid (CANON_SONIC §6).
- Conversion **adds no precision**. A surplus known to two decimals is not known to four because cents has finer graduation. See F-32.7.

The Specification

Numbered, checkable requirements. Any document, chapter, tool or report in this corpus that crosses registers must satisfy all of them.

CK-1. Every logarithmic quantity is published with its **base** (nats, bits, octaves, cents) in the same line or column header. A bare number is not a measurement.

CK-2. Every logarithmic quantity that is not a ratio between two reported values is published with its **reference** — `I_0`, `f_ref`, or the ground essence `k` — stated explicitly.

CK-3. Every published quantity carries its **status badge**: PROVEN, DEFINITION, MEASURED, POSTULATE, or CHOSEN NORMALISATION. Conversions preserve the badge.

CK-4. The surplus-to-pitch conversion is published as `octaves = 2.885390 x F` **with the note that the factor 2 is the exponent of the essence-law postulate**, not an independent constant.

CK-5. Any published essence frequency carries, in the same caption or table note: *this is an essence frequency, not a physical resonance.*

CK-6. Any essence frequency published as an audible value states the **band** and the integer **j** of the transposition `f x 2^j` (CANON_SONIC §5), and a set of values whose intervals are the point shares one `j`.

CK-7. Any surplus figure from the deployed functional is published as a **lower bound**, with "at least" or >=.

CK-8. Stevens' exponents are published with the **procedure**: method (estimation or production), stimulus range, instruction. An exponent without a procedure is not a measurement.

CK-9. Weber fractions are never asserted for a modality unless measured in-house or cited. Parametric tables over `c` are the permitted substitute and are labelled parametric.

CK-10. Distortion is reported as `Delta = |ln r|` in a stated register, **never as the raw ratio `r` alone**, which is asymmetric between over- and under-framing.

CK-11. The Moment equation is never published without its badge, nor without the note that `D` and `P` must share units.

CK-12. "Maximum moment = `ln N`" is published with its derivation chain — Paper D Cor 2.2 plus the chosen normalisation — both links, every time.

CK-13. Precision after conversion never exceeds precision before it. Report to the source's **significant figures**, not to the target register's natural resolution.

CK-14. Any table crossing registers is **computed**, not transcribed, and is reproducible from the constants in §32.3.

CK-15. Where a conversion depends on the essence exponent `s = 2`, the dependence is stated and the sensitivity table (§32.6) is available.

CK-16. No conversion between a ratio-scaled and an interval-scaled quantity without demoting the result and saying so. Specifically: **Fechner sensations may be subtracted but not divided**.

CK-17. Every formula entering a specification carries its full seven-gate FVP table in print, including gates it fails.

CK-18. Failed and unattempted gates print as **FAIL** or **NOT YET ATTEMPTED** — never omitted, softened, or deferred to "future work" without naming the specific test.

CK-19. No intervention derived from these conversions removes a person's ability to decline it (CANON §5.4). With special force for anything delivered as tone: **audible, disclosed, refusable, with no subliminal content of any kind** (CANON_SONIC §7).

CK-20. Any claim that two registers are "the same thing" is prohibited. They are inter-convertible measures of quantities sharing a logarithmic structure. That is the claim, and it is enough.

The Failure Modes

F-32.1 . THE UNSTATED BASE. *Signature:* a surplus, sensation or pitch figure published as a bare number; readers assume bits where nats were meant. *Damage:* a silent factor of 1.442695 or 0.693147 through every derived quantity. In headroom calculations this is the difference between `h > 0` and `h < 0`. *Detection:* audit any figure whose ratio to a companion figure sits near 1.4427 or 0.6931. *Fix:* CK-1. Unit in the column header, never removed.

F-32.2 . THE FREE LUNCH CONVERSION. *Signature:* a modest surplus gain displayed in Hertz, where it looks like a transformation. "We moved from 2 Hz to 20 Hz" is `N_eff` from 2.03 to 6.43 — 1.15 nats, real but a third of what the Hertz figure implies. *Damage:* over-claiming, then loss of credibility when the figure is recomputed. *Detection:* recompute in nats. If the nat figure is unimpressive, the Hertz figure was theatre. *Fix:* nats or bits in the argument; Hertz only in the sonification.

F-32.3 . THE PHYSICAL RESONANCE SLIP. *Signature:* "the system resonates at", "tuned to the frequency of", "vibrates at", applied to an essence frequency. *Damage:* the framework is correctly reclassified as pseudoscience and the genuine results are discarded with the error. *Detection:*

grep for “resonat”, “vibrat”, “frequency of the”. Every hit must concern an actual physical system or carry the disclaimer. *Fix:* CK-5, without exception and without weariness.

F-32.4 . THE BORROWED EXPONENT. *Signature:* a Stevens exponent measured by one procedure, imported into a specification using a different range, instruction, or modality sub-type. *Damage:* the key is wrong by an unknown amount, and because the error is in an exponent it compounds with the range. *Detection:* ask what procedure produced it. “The literature” means borrowed. *Fix:* CK-8. Measure in-house, or mark nominal and carry the uncertainty forward.

F-32.5 . THE LINEAR PROMISE. *Signature:* a commitment to “twice the brightness” or “twice the loudness”, costed as though the physical quantity doubles. *Damage:* the real requirement is `2^(1/a)` — for brightness, 8.17x. Budgets built on the linear assumption are short by a factor of four. *Detection:* any specification whose perceptual target and physical budget are the same multiple. *Fix:* the Lens Theorem (Chapter 36). **Never promise a linear improvement in a quantity the recipient perceives logarithmically.**

F-32.6 . THE UNBOUNDED MOMENT. *Signature:* a claim that experience can be arbitrarily enriched by intensity, novelty or attention, without reference to the perceiver’s category count. *Damage:* effort spent past the point where `ln N` binds, with no explanation available for the vanishing returns. *Detection:* compute `F_int / ln N`. Above about 0.9 the ceiling binds and geometry offers nothing. *Fix:* raise `N`, and accept that the gain is logarithmic. Distinguish from the *raw* `M`, which is unbounded as `P -> 0` and whose divergence is pathology. **Hallucination is unbounded moment; it is not a rich moment.**

F-32.7 . THE CENTS ILLUSION. *Signature:* a surplus estimated to two significant figures, reported as 4747.67 cents. *Damage:* six figures of apparent precision on two figures of information; readers treat the difference between 4747 and 4750 cents as meaningful. It is the fourth decimal of `u`. *Detection:* compare significant figures in source and conversion. *Fix:* CK-13. Convert, then round to the source’s precision. Report `u` and `N` if you want the reader to know what was measured.

F-32.8 . THE MISSING ZERO. *Signature:* division of two quantities lacking a true zero — most often ordinal ratings treated as `D` and `P`. *Damage:* `r` is arbitrary, `|ln r|` is arbitrary, and every downstream conversion is arbitrary while looking rigorous. *Detection:* FVP-2. Ask what zero would mean for each term and whether the instrument can produce it. *Fix:* build a ratio scale (Chapter 53), or demote the formula to a heuristic ordering **in print**.

F-32.9 . THE STATUS UPGRADE. *Signature:* a postulate stated in nats in one chapter, cited in octaves in the next, treated as established by the third. *Damage:* a design decision acquires the authority of a theorem. *Detection:* trace any load-bearing constant to its first appearance and check the badge there. *Fix:* CK-3. The badge travels with the number, not with the sentence.

Worked Example

A metropolitan transit review board, carried end to end in four registers.

Survey

A municipal transit authority convenes a review board on a proposed bus-priority corridor. The board has **N = 8 declared professional categories**: traffic engineering, urban planning, transit operations, accessibility, municipal finance, local business representation, environmental assessment, and legal/regulatory.

Non-parallelism u is estimated by the disagreement-rate protocol of Chapter 7: across 40 recorded recommendations, the rate at which two randomly drawn members' initial positions diverge, corrected for categorical structure. **Measured u = 0.42.** This is an in-example figure produced by a stated procedure, not a citation.

Storage parameters, measured over four quarters by the protocols of Chapters 3-5:

```
eta   = 0.35     fraction of board surplus that becomes durable policy structure
delta = 0.08       quarterly decay of adopted structure without maintenance
Q     = 4.0      current stored order (nat-equivalent, on the board's own scale)
C     = 0.30         quarterly coordination cost in the same units
```

Step 1 — Effective count and surplus

```
N_eff = 1 + (N - 1) u = 1 + 7 * 0.42 = 3.94
F     = ln 3.94       = 1.371181 nats
```

Eight declared categories are worth 3.94 genuinely different ones. More than half the board's declared variety is redundant, which is the ordinary finding and not a scandal.

Step 2 — The four registers

register	value	reading
surplus	**>= 1.371181 nats**	lower bound on this board's per-session information
information	>= 1.978196 bits	just under two bits: the board resolves fewer than four states
essence pitch	3.956391 octaves above ground	
cents	4747.7 cents	rounded to 4750 per CK-13, since u is known to two figures
essence frequency	f = k . N_eff² = 7.504508 Hz	**essence frequency, not a physical resonance**
audible placement	120.07 Hz, band [110, 220) Hz, **j = 4**	

The bit figure is the one to put in front of the board. **Two bits means this body, meeting as it currently meets, distinguishes fewer than four genuinely different states of the world per session.**

Step 3 — The ceiling and the gap

```
ceiling   ln N = ln 8 = 2.079442 nats = 3.000000 bits exactly = 6 octaves exactly
 achieved  F/ln N = 1.371181 / 2.079442 = 0.6594
 gap       ln N − F = 0.708261 nats = 2.0436 octaves = 2452 cents
```

The board is at **66% of its own ceiling.** The ceiling frequency is `k . 8² = 30.939247 Hz`, and the ratio of ceiling to achieved frequency is `(8/3.94)² = 4.1228` — **the board is two full octaves below its own maximum**, the same fact as "34% of the ceiling unused", said in the register that makes it audible.

Step 4 — The storage arithmetic

```
headroom      h  = eta∗ln N − delta∗Q − C
                 = 0.35 ∗ 2.079442 − 0.08 ∗ 4.0 − 0.30
                 = 0.727805 − 0.320000 − 0.300000
               = +0.107805       h > 0: growth is available in principle

actual step      = eta∗F − delta∗Q − C
                 = 0.35 ∗ 1.371181 − 0.320000 − 0.300000
                 = 0.479913 − 0.620000
                 = −0.140087        Q is DECLINING
```

Here is the diagnostic the conversion keys make visible: positive headroom, negative actual step. The ceiling permits growth; the geometry does not deliver it. The authority is losing adopted structure at 0.14 nat-equivalents per quarter while believing itself functional, because nothing in its minutes records a quantity that would show the decline.

Step 5 — What is required

By Proposition 4.1 (Paper E), a non-decreasing step requires

```
S >= (delta∗Q + C) / eta = (0.320 + 0.300) / 0.35 = 1.771429 nats

Feasible?  1.771429 <= ln 8 = 2.079442     YES, with 0.308 nats to spare
```

Solving back for the geometry:

```
N_eff required = e^1.771429 = 5.879246
u    required  = (5.879246 − 1) / 7 = 0.697035
```

u must rise from 0.42 to 0.70 — a climb of `2 x (1.771429 − 1.371181) / ln 2 = 1.154871` octaves, or 1385.85 cents (report: 1390 cents), with the essence frequency rising from 7.504508 Hz to 16.709870 Hz, a factor of 2.2266.

Step 6 — Specification and prediction

The binding constraint is u, not N. Adding a ninth category raises the ceiling from `ln 8 = 2.0794` to `ln 9 = 2.1972` nats — **0.1178 nats, one third of a bit, for the cost of an entire additional seat** — while doing nothing about the redundancy among the eight already present. Chapter 9's levers apply: rotation across lineages, a standing adversarial role, a member from an unrelated

trade, and an explicit prohibition on the pre-meeting alignment call that currently drives u toward zero before the session begins.

Predicted outcome, stated so it can fail: if u rises to 0.70 or above and `eta`, `delta`, `C` hold, the quarterly step turns non-negative within two quarters and `Q` stabilises near `Q_inf = (eta.F − C)/delta = (0.35 x 1.771429 − 0.300)/0.08 = 3.9999`, i.e. approximately its current level — **the intervention stops the decline; it does not build.** Building requires raising `eta` or cutting `C`, which is Chapter 26's lever ranking and is a separate decision.

Falsification: if u measures above 0.70 for two consecutive quarters and `Q` still declines, then either `eta` is below 0.35, or `C` is above 0.30, or the linear storage model is wrong for this body. All three are checkable; the third is the one that would matter. **A framework that cannot fail cannot be trusted, and this one names exactly how it fails.**

Step 7 — What the sonification is for

If this board's trajectory is sonified, both values take **the same** transposition, `j = 4`: 120.07 Hz now, 267.36 Hz at the target, band [110, 440) Hz. The tone rises by 1.1549 octaves and the change is immediately audible where a table of nats is not.

The shared `j` is not a detail. Placing each value in a narrow band with its own `j` — 120.07 Hz at `j = 4` and 133.68 Hz at `j = 3` — preserves each pitch class but destroys the interval between them, collapsing 1.15 octaves to 0.15. Octave transposition preserves pitch class exactly and preserves intervals only when the transposition is shared. **Publish the band and the `j`, and use one `j` across any set of values whose intervals are the point** (CK-6).

That said, sonification is a display, and displays are for perception.

It is not evidence. Nothing about the transit authority vibrates. The tone is an essence frequency computed from a measured geometry by a postulated square law with a chosen reference constant, and every one of those four qualifications belongs in the caption. The board should hear it, understand exactly what it is, and be free to ignore it.

The keys are in hand. Chapter 33 asks the question the honesty clauses force: **what makes an ordering of categories defensible enough to put a frequency on it?**

Chapter 33 . Harmonic Lattices and Essence Sets

Chapter 32 established that a surplus can be written as a pitch by an exact constant. It also established the condition under which that is worth doing, and then deferred it. The condition is this: **the essence law consumes an ordinal index, and an arbitrary ordinal produces an arbitrary frequency.** The law will happily convert nonsense. The whole burden of applying it therefore falls on the construction of the index — on the essence set.

This chapter discharges that burden. It does two things. First it works out the **structure** of the lattice the square law generates: what intervals exist, how they fold, how many distinct pitch classes an N-essence set can possibly have, and which pairs are consonant. All of that is exact arithmetic and all of it is computed below. Second it states the **requirements** an ordering must meet before a frequency may be attached to it, and demonstrates them on a real, published, externally-ordered category set.

The structural results are unusually clean, and one of them is a genuine constraint rather than a convenience: **an essence set of N members contains exactly ceil(N/2) distinct pitch classes, no matter what the essences are.** Half your vocabulary is octave duplication before you begin. That is a hard limit on the harmonic vocabulary available to any essence set, and it is not a stylistic choice.

The Law

33.1 The essence law and its status

From *The Effective Method of Sonic Chemistry* (Curzi, 17 January 2022), collapsed to its exact form (CANON_SONIC §1):

```
f(n) = k * n^2                              [POSTULATE about essence structure]

k = (9/8)^2 / phi^2 = (81/64)/(phi + 1)          [CHOSEN reference constant]
   = 0.4834257329884143 Hz

 n = the ESSENCE INDEX: the ordinal position of an essence type within its
    discipline's ordered category set               [REQUIRES JUSTIFICATION]
```

Three statuses in three lines, and none of them is “proven”. The square law is a postulate about how essence types relate. k is a design decision — a Pythagorean ditone divided by the golden section — and it is not measured. The index is an input that the law does not supply and cannot check.

The output is an ESSENCE frequency. It is the resonance of the essence, not a physical resonance. No claim is made, here or anywhere, that any material, molecule, element, room, institution, person, or category physically vibrates at these values. This clause is restated at every point in this chapter where a frequency appears, because the temptation to read it otherwise is the single most dangerous failure mode in the material.

Verification of the collapse, at 17 significant figures:

```
 phi                = 1.618033988749895
```

```
phi^2              = 2.618033988749895      = phi + 1, as required
(9/8)^2 / phi^2    = 0.48342573298841435
(81/64) / (phi+1)  = 0.48342573298841435    identical
```

33.2 Result One — every interval is a ratio of squares, and k cancels

```
f(m) / f(n) = (k m^2) / (k n^2) = m^2 / n^2
```

The golden ratio and the 9/8 set the absolute reference pitch and nothing else. Every relationship in the system is a ratio of squared integers. This is the most important structural fact about the lattice, because it means the chosen constant k — the one part of the law with no justification beyond the author's preference — has **no influence whatever on any interval, consonance, pitch class, or harmonic relation.** It sets where the instrument sits and not how it is tuned.

Verified against the examples recorded in CANON_SONIC §2, computed here independently:

```
index ratio   essence ratio   cents      description
2 : 1          4 / 1          2400.00    two octaves exactly
3 : 2        9 / 4          1403.91    octave + 203.91 cents (a whole tone)
4 : 3         16 / 9           996.09    a minor seventh
5 : 4         25 / 16          772.63    near a minor sixth
3 : 1          9 / 1          3803.91    three octaves + 203.91 cents
```

All five agree with the canon. The 203.91 cents appearing twice is not a coincidence: it is `1200 * log2(9/8)`, the Pythagorean whole tone, and it is the pitch class of essence index 3.

33.3 Result Two — pitch class is the odd part of the index

Because `(2n)^2 = 4n^2` is exactly two octaves above n^2, multiplying an index by two transposes by a whole number of octaves and changes nothing else.

Write any index in the form `n = 2^a * m` with `m` odd. Then

```
f(n) = k * (2^a * m)^2 = 4^a * (k m^2) = k m^2 transposed up by 2a octaves
```

Two essences share a pitch class if and only if their indices have the same odd part. The pitch class is fully determined by `m`, and its position within the octave is `m^2` reduced by powers of two.

THE CLASS-COUNT THEOREM. *An essence set indexed 1..N contains exactly `ceil(N/2)` distinct pitch classes.*

Proof. The pitch class of index n is its odd part. For any odd `m` with `1 <= m <= N`, the index `m` itself lies in `1..N` and has odd part `m`, so every odd number up to N is realised. Conversely the odd part of any `n <= N` is itself at most n, hence at most N, and is odd. So the set of realised classes is exactly the set of odd integers in `[1, N]`, which has `ceil(N/2)` elements. []

Verified by direct enumeration for `N = 1..24` (left column N, centre the counted number of distinct odd parts, right `ceil(N/2)`):

```
  N :  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
cnt :  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12
```

```
ceil :  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12
       all 24 agree
```

Twelve essences yield only six pitch classes. That is the available harmonic vocabulary, and it is a hard structural limit.

33.4 The three requirements on an essence set

Before any of the above may be applied to a real domain, the domain's categories must be turned into an index. CANON_SONIC §6 requires "an explicit, justified ordinal", and Paper B's validity conditions on state encoding and category assignment govern. Made operational, this is three requirements:

R1 — EXHAUSTIVE AND DISJOINT. The categories partition the domain. Every state of the domain falls in exactly one. If two categories overlap, an entity has two indices and therefore two frequencies, and the lattice is not well-defined.

R2 — AN ORDERING FROM OUTSIDE THE MUSIC. The ordinal must be fixed by a property of the domain that was established before, and independently of, any intention to sonify it. **The ordering may not be chosen to make the result sound good.** This is the essence-set analogue of preregistration and it is the requirement that is most often violated.

R3 — REPRODUCIBLE ASSIGNMENT. A second competent analyst, given the domain and the stated ordering rule, produces the same index for the same category. This is FVP-7 applied to the input rather than the output, and it is checkable by inter-rater agreement.

A set failing any of R1-R3 may not be sonified. It may still be described, tabulated, and reasoned about — it simply does not get a frequency.

33.5 The scale-promotion problem, badged

Here is the sharpest honest difficulty in this material, and it must be stated in the open.

The essence law takes an ORDINAL and returns a RATIO-SCALE quantity. The index is a rank. The frequency is a physical magnitude with a true zero, on which multiplication and division are meaningful. **The law therefore performs a scale promotion — it manufactures ratio-scale structure out of ordinal input — and by FVP-2 that is not a permissible operation on data alone.**

It is permissible as a **model**, and that is what it is. The square law is a postulate that *asserts* a specific quantitative spacing between ranks. The assertion may be right or wrong, and it is the thing that would have to be tested. **What the text may never do is present the resulting frequencies as though the ratio-scale structure were measured. It was postulated.**

The practical form of this rule: **the intervals are only as meaningful as the spacing postulate. The ordering is data; the spacing is the model.**

33.6 The honesty clauses, restated

- **The output is an essence frequency, not a physical resonance.** Nothing vibrates.
- **k is a chosen reference constant**, not a measured one, and it cancels from every interval.

- **The square law is a postulate about essence structure.** Its formal coincidence with the eigenvalue spacing of a one-dimensional infinite square well, where E_n is proportional to n^2, is a genuine mathematical correspondence and may be noted — and coincidence of form is not evidence of physical mechanism, which must be said in the same breath.
- **The essence index requires a defensible ordering** (R1-R3); an arbitrary ordinal yields arbitrary frequencies.
- **The consonance indices used below are CHOSEN measures**, defined in §33.10, not measured psychoacoustic facts about human listeners.
- **The surplus bridge inherits all of Paper A-E's honesty clauses**, including that surplus is a lower bound on ensemble information and that no thermodynamic or physical claim follows.
- **No subliminal content of any kind.** Anything delivered as tone is audible, disclosed, and refusable (CANON_SONIC §7, CANON §5.4).
- **Recompute; never copy a published value from the 2022 source.** Chapter 29 §29.4 established that `f = k n^2` reproduces all 118 published element values to a worst relative deviation of **9.0e-15**, and §29.10 that the compound rule `f = k^2 G^4` — G the count-weighted geometric mean of the constituent indices — matches **37 of the 44** published composites to better than one part in 10^12. **The seven that disagree are data-entry errors in the source, not evidence of a different law, and the O2 row proves it: for a single-element species G is the element index whatever the atom count, so no atom count whatever can produce the published figure.** Two transcription faults travel with the citation: the source's prose instruction `1+5^.5/2` is missing its parentheses and must be read as `(1 + sqrt(5))/2`, and the spreadsheet stored `phi` to fifteen significant figures, which accounts for the whole of the residual bias in the published table. Every frequency in this chapter is recomputed from the corrected constants; **cite the errata whenever the source is cited.**

The Engineering

33.7 The lattice, computed for N = 24

The complete essence set for indices 1 to 24. Columns: raw essence frequency `k n^2`; the factorisation `n = 2^a . m`; the pitch class `m`; the position of n above the ground essence in cents (`2400 . log2 n`); and the pitch class's position within the octave (that value modulo 1200).

n	f = k n^2 Hz	n = 2^a*m	class m	cents above ground	class cents
1	0.483426	2^0*1	1	0.0000	0.0000
2	1.933703	2^1*1	1	2400.0000	0.0000
3	4.350832	2^0*3	3	3803.9100	203.9100
4	7.734812	2^2*1	1	4800.0000	0.0000
5	12.085643	2^0*5	5	5572.6274	772.6274
6	17.403326	2^1*3	3	6203.9100	203.9100
7	23.687861	2^0*7	7	6737.6518	737.6518
8	30.939247	2^3*1	1	7200.0000	0.0000
9	39.157484	2^0*9	9	7607.8200	407.8200
10	48.342573	2^1*5	5	7972.6274	772.6274
11	58.494514	2^0*11	11	8302.6359	1102.6359
12	69.613306	2^2*3	3	8603.9100	203.9100

13	81.698949	2^0*13	13	8881.0553	481.0553
14	94.751444	2^1*7	7	9137.6518	737.6518
15	108.770790	2^0*15	15	9376.5374	976.5374
16	123.756988	2^4*1	1	9600.0000	0.0000
17	139.710037	2^0*17	17	9809.9108	209.9108
18	156.629937	2^1*9	9	10007.8200	407.8200
19	174.516690	2^0*19	19	10195.0260	595.0260
20	193.370293	2^2*5	5	10372.6274	772.6274
21	213.190748	2^0*21	21	10541.5618	941.5618
22	233.978055	2^1*11	11	10702.6359	1102.6359
23	255.732213	2^0*23	23	10856.5487	56.5487
24	278.453222	2^3*3	3	11003.9100	203.9100

Every essence frequency in that table is an essence frequency and not a physical resonance.

33.8 The pitch classes and their octave positions

Folding the table above by odd part gives the complete harmonic vocabulary of a 24-essence set: twelve classes, as the theorem requires.

class m	member indices <= 24	m^2 reduced to [1,2)	cents	Hz in [k, 2k)
1	1, 2, 4, 8, 16	1/1	0.000	0.483426
23	23	529/512	56.549	0.499477
3	3, 6, 12, 24	9/8	203.910	0.543854
17	17	289/256	209.911	0.545742
9	9, 18	81/64	407.820	0.611836
13	13	169/128	481.055	0.638273
19	19	361/256	595.026	0.681706
7	7, 14	49/32	737.652	0.740246
5	5, 10, 20	25/16	772.627	0.755353
21	21	441/256	941.562	0.832776
15	15	225/128	976.537	0.849772
11	11, 22	121/64	1102.636	0.913977

Listed in pitch order, and the ordering is worth noting: **the pitch classes do not appear in index order.** Class 23 sits 56.5 cents above the ground, below class 3 at 203.9 cents, even though its index is nearly eight times larger. Squaring and folding scrambles the ordinal completely. **The lattice preserves the index in absolute pitch and destroys it in pitch class.** An analyst who reads pitch-class order as importance order has misread the arithmetic.

Two structural features of that list deserve engineering attention.

The vocabulary is unevenly distributed. The gaps between consecutive classes, in cents:

```
  0.000 -> 56.549    gap  56.549     (class 1 to class 23)
 56.549 -> 203.910   gap 147.361
203.910 -> 209.911   gap   6.001     (class 3 to class 17)
209.911 -> 407.820   gap 197.909
407.820 -> 481.055   gap  73.235
481.055 -> 595.026   gap 113.971
595.026 -> 737.652   gap 142.626
```

```
 737.652 -> 772.627   gap  34.976     (class 7 to class 5)
 772.627 -> 941.562   gap 168.934
 941.562 -> 976.537   gap  34.976     (class 21 to class 15)
 976.537 -> 1102.636  gap 126.098
1102.636 -> 1200.000  gap  97.364     (wrapping to class 1)
```

Four pairs sit inside 60 cents of each other, which is close enough that as simultaneous tones they will beat rather than form an interval:

```
classes  3 and 17 :   6.00 cents apart   ( 9/8  vs 289/256 )
classes  7 and  5 :  34.98 cents apart   ( 49/32 vs 25/16 )
classes 21 and 15 :  34.98 cents apart   ( 441/256 vs 225/128 )
classes  1 and 23 :  56.55 cents apart   ( 1/1  vs 529/512 )
```

The 6.00-cent gap between classes 3 and 17 is the one to watch. **Essence index 3 and essence index 17 are, for practical purposes, the same pitch class, and nothing in the ordering of the domain caused that — it is an artefact of 9/8 and 289/256 being nearly equal.** If a domain's third and seventeenth categories are conceptually remote, the lattice will nevertheless render them as a near-unison, and any listener will hear them as one blurred tone. The 34.98-cent coincidence between classes 5 and 7 (and identically between 15 and 21, since `21/15 = 7/5`) is the same phenomenon at a coarser scale.

The design consequence is direct: essence sets past a certain size begin to produce near-unison collisions that the ordering did not intend and cannot control. This is an argument for small essence sets that has nothing to do with taste.

And the size is exact, which is better than "about twelve". *This sharpens the statement above, which was stated approximately in earlier drafts of this chapter.* Computing the folded lattice's minimum gap in exact rational arithmetic for every set size gives a step function with two values and one cliff:

```
 N = 3 .. 6      minimum gap  9/8       = 203.910 cents     no collision
N = 7 .. 16     minimum gap  50/49     =  34.976 cents     classes 7 and 5
N = 17 .. 40    minimum gap  289/288   =   6.001 cents     classes 17 and 3
N = 41 and up   minimum gap  1682/1681 =   1.030 cents     classes 41 and 29
```

The third step is recorded for completeness only: `N = 41` lies far beyond the 25-essence ceiling that §33.13 derives for a shared transposition, so no renderable set reaches it.

The first coincidence arrives at `N = 7`, when class 7 lands 34.976 cents from class 5. The lattice then holds that gap, unchanged, for ten consecutive set sizes — and collapses by a factor of 8 in achievable rate (5.828 in cents; see Ch 29 Law 2) the moment index 17 enters, because `17^2 = 289` sits one part in 288 above `9/8 = 288/256`. Nothing in the domain causes either event; both are arithmetic.

The cliff is not merely a narrower gap, it is a change of regime, and Chapter 71 prices it: by the surplus-resolution bound, telling two tones apart requires `1/(8 pi dF)` cycles, so a 34.976-cent separation needs 3.94 cycles and a 6.001-cent separation needs 22.96. **A set of sixteen and a set of seventeen require rendering rates differing by a factor of exactly 8, and the boundary between them is exact.** Design at 16 or below if you can; if you cannot, budget the tone duration before you budget anything else. (Chapter 71 states this boundary at `N = 19` to `N = 20`; the value is corrected here to `N = 16` to `N = 17`, which is what the exact lattice returns and what

§33.7's own table shows — index 17 is present in a 24-set and its 6.001-cent gap with class 3 is listed above.)

FVP . The Essence Law

`f(n) = k . n^2` — **[POSTULATE about essence structure]**

Gate	Test	Result
FVP-1 DIMENSIONS	`n` is a dimensionless index; `k` carries Hz; `f` is in Hz. Both sides agree. `k` must be published with its units or the equation is incomplete.	**PASS**
FVP-2 SCALE TYPE	**The crux.** `n` is ORDINAL. `f` is RATIO-scaled. The law promotes an ordinal to a ratio scale by postulating the spacing. Permissible as a model, impermissible as an inference from data. **Must be badged in print wherever intervals are interpreted.**	**PROMOTION BY POSTULATE — badged, not passed**
FVP-3 LIMITING CASES	`n = 0` gives `f = 0`: there is no zeroth essence, and the law says so by returning silence. `n = 1` gives `f = k`, the ground essence. `n -> infinity` diverges quadratically, leaving the audible band quickly — at `n = 24`, `f` is already 278 Hz and at `n = 200` it is 19.3 kHz. Transposition (§33.13) is the required remedy.	**PASS, with a declared working range**
FVP-4 FIXED POINTS	`f(n) = n` has solutions at `n = 0` and `n = 1/k = 2.0685...`, neither of which is meaningful, since index and frequency are not the same kind of quantity. **The meaningful invariant is `f(1) = k`: the ground essence is the fixed reference and every other value is relative to it.**	**PASS — invariant is the ground, not a self-map**

Gate	Test	Result
FVP-5 IDENTIFIABILITY	k and the exponent 2 are separately recoverable from any two published values: two points determine both slope and intercept in log-log coordinates. **The exponent was recovered this way from the 2022 source document and matched a pure square to 14 significant figures**, which is how the collapse in CANON_SONIC §1 was established.	**PASS — and was actually performed**
FVP-6 FALSIFICATION	As arithmetic, not falsifiable. As a claim about essence structure, falsifiable only once "essence relationship" has an independent operationalisation with its own measurable consequence. **No such operationalisation currently exists, and the honest statement is that the postulate is not yet testable.** Naming what would be needed: a domain where the relation between category types has an independently measured magnitude, against which the postulated `m^2/n^2` spacing could be checked.	**NOT CURRENTLY FALSIFIABLE — stated plainly**
FVP-7 REPRODUCIBILITY	The arithmetic reproduces exactly. **The INDEX does not reproduce automatically** — it depends on the ordering, and the ordering is where reproducibility must be established, by R3 inter-rater agreement.	**CONDITIONAL — the gate transfers to the ordering**

33.9 The interval lattice in cents

The interval between essence `i` and essence `j` is `2400 . log2(j/i)` cents. The full matrix for indices 1..12, in cents, upper triangle (the lower triangle is the negative):

	j=2	3	4	5	6	7	8	9	10	11	12
i=1	2400	3804	4800	5573	6204	6738	7200	7608	7973	8303	8604
i=2	–	1404	2400	3173	3804	4338	4800	5208	5573	5903	6204
i=3	–	–	996	1769	2400	2934	3396	3804	4169	4499	4800
i=4	–	–	–	773	1404	1938	2400	2808	3173	3503	3804
i=5	–	–	–	–	631	1165	1627	2035	2400	2730	3031
i=6	–	–	–	–	–	534	996	1404	1769	2099	2400
i=7	–	–	–	–	–	–	462	870	1235	1565	1866
i=8	–	–	–	–	–	–	–	408	773	1103	1404
i=9	–	–	–	–	–	–	–	–	365	695	996
i=10	–	–	–	–	–	–	–	–	–	330	631
i=11	–	–	–	–	–	–	–	–	–	–	301

Read the diagonal band — the intervals between adjacent indices — and a pattern appears that is worth stating as a result in its own right.

adjacent pair	raw interval (cents)
1 – 2	2400.00
2 – 3	1403.91
3 – 4	996.09
4 – 5	772.63
5 – 6	631.28
6 – 7	533.74
7 – 8	462.35
8 – 9	407.82
9 – 10	364.81
10 – 11	330.01
11 – 12	301.27
12 – 13	277.15
15 – 16	223.46
19 – 20	177.60
23 – 24	147.36

Neighbouring essences crowd together as the index rises. The gap between the first two essences is two octaves; the gap between the twenty-third and twenty-fourth is 147 cents, less than a whole tone and a quarter. This is the logarithm again: `2400 . log2((n+1)/n)` falls like `2400/(n ln 2)`. In an essence set of 24, the top of the set is harmonically congested and the bottom is spread across octaves.

Consequence for set design: the resolution of the lattice is not uniform, and low indices carry far more pitch distance than high indices. If the domain's ordering is such that its most important distinctions land at high index, the lattice will render them as small intervals. That is a reason to think hard about which end of the ordering the important categories occupy — and, critically, **it is not a licence to reorder the set to improve the result**, which R2 forbids.

FVP . The Interval Law

```
f(m)/f(n) = m^2 / n^2
```

Gate	Test	Result
FVP-1 DIMENSIONS	A ratio of two frequencies: dimensionless. `k` cancels identically, so the result carries no dependence on the chosen reference. Both sides dimensionless.	**PASS**
FVP-2 SCALE TYPE	The ratio is well-formed on a ratio scale, which `f` is by construction. **But the ratio-scale status of `f` was itself manufactured by the postulate (§33.5), so the interval inherits that status, not a stronger one.**	**PASS given the postulate, badged**
FVP-3 LIMITING CASES	`m = n` gives 1: unison, an essence with itself. `n -> 0` diverges, and there is no zeroth essence, so the divergence is outside the domain. `m/n -> infinity` gives unbounded interval, correctly: distant indices are distant pitches.	**PASS**
FVP-4 FIXED POINTS	The ratio equals 1 iff `m = n`; it equals a whole number of octaves iff `m/n` is a power of two. **The invariant of the folding operation is the odd part, which is Result Two.**	**PASS — fixed point is the class**
FVP-5 IDENTIFIABILITY	The interval determines `m/n` uniquely (take the square root of the ratio). It does not determine `m` and `n` separately — `6:4` and `3:2` give the same interval. **The lattice is identifiable up to the reduced index ratio and no further**, which is exactly why §33.10's consonance depends on the reduced ratio alone.	**PASS with a named degeneracy that turns out to be structural**

Gate	Test	Result
FVP-6 FALSIFICATION	Follows the essence law: falsifiable only when the underlying spacing postulate becomes testable. The interval law adds no new empirical content beyond it.	**INHERITS THE ESSENCE LAW'S STATUS**
FVP-7 REPRODUCIBILITY	Exact arithmetic, reproducible anywhere. Verified independently in this chapter against the five examples in CANON_SONIC §2 — all five agree.	**PASS**

33.10 Consonance as a function of the index ratio

Which pairs of essences sound related, and which sound remote? The question has a precise answer in this lattice, and the answer depends **only on the reduced ratio of the indices**.

Let `i/j` reduce to `p/q` with `gcd(p, q) = 1`. Then the essence interval is `p^2/q^2`, and since `gcd(p^2, q^2) = 1`, that fraction is already in lowest terms. Everything about the pair's harmonic relationship is therefore a function of `p` and `q` alone.

Two computable indices are used here. **Both are CHOSEN measures with a defined arithmetic, not measured facts about human hearing.** They are badged as such, and where they disagree the text says so.

```
    TENNEY HARMONIC DISTANCE       HD(p/q) = log2(p * q)
     A defined measure of the complexity of a frequency ratio. Larger = more remote.

   EULER'S GRADUS SUAVITATIS    G(n) = 1 + sum over prime factors of e_i * (prime_i -
1)
     Euler, Tentamen novae theoriae musicae, 1739. G(p/q) = G(p*q). Larger = less smooth.
```

THE DOUBLING THEOREM. *For any index pair reducing to `p/q`, the essence interval `p^2/q^2` has exactly twice the Tenney harmonic distance of the ratio `p/q`, and Euler gradus `2G - 1`.*

Proof. `HD(p^2/q^2) = log2(p^2 q^2) = 2 log2(pq) = 2 HD(p/q)`. For the gradus: squaring doubles every exponent in the prime factorisation of `pq`, so the sum `Σ e_i (prime_i - 1)` doubles, and `G(p^2 q^2) = 1 + 2 Σ e_i (prime_i - 1) = 2(1 + Σ e_i (prime_i - 1)) - 1 = 2G(pq) - 1`. []

Verified exhaustively over all pairs `1 <= j < i <= 12` — 66 pairs, both identities exact in every case. A sample:

```
  i   j    p/q    HD(p/q)   HD(p^2/q^2)   2*HD    G(p/q)   G(ess)   2G-1    cents
  2   1    2/1     1.0000      2.0000    2.0000     2        3       3     2400.00
  3   2    3/2     2.5850      5.1699    5.1699     4        7       7     1403.91
```

4	3	4/3	3.5850	7.1699	7.1699	5	9	9	996.09
5	4	5/4	4.3219	8.6439	8.6439	7	13	13	772.63
6	5	6/5	4.9069	9.8138	9.8138	8	15	15	631.28
7	6	7/6	5.3923	10.7846	10.7846	10	19	19	533.74
8	7	8/7	5.8074	11.6147	11.6147	10	19	19	462.35
8	3	8/3	4.5850	9.1699	9.1699	6	11	11	3396.09
7	5	7/5	5.1293	10.2586	10.2586	11	21	21	1165.02

This is the deepest structural result in the chapter. The square law systematically doubles harmonic distance. Every essence pair is exactly twice as harmonically remote as the corresponding integer-ratio pair would be in ordinary just intonation. 3:2 — the perfect fifth, the most consonant interval after the octave — becomes 9:4, an octave plus a whole tone, at double the harmonic distance. 5:4, the major third, becomes 25:16.

Consonance is scarce in the essence lattice, and it is scarce by a factor of exactly two.

33.11 The consonance ranking, and what it reveals

Ranking all distinct index ratios among indices 1..12 by essence harmonic distance:

HD_ess	Gradus	index ratio	essence ratio	reduced	cents
0.0000	1	1:1	1/1	1	0.00
2.0000	3	2:1	4/1	1	0.00
3.1699	5	3:1	9/1	9/8	203.91
4.0000	5	4:1	16/1	1	0.00
4.6439	9	5:1	25/1	25/16	772.63
5.1699	7	3:2	9/4	9/8	203.91
5.6147	13	7:1	49/1	49/32	737.65
6.0000	7	8:1	64/1	1	0.00
6.3399	9	9:1	81/1	81/64	407.82
6.6439	11	5:2	25/4	25/16	772.63
6.9189	21	11:1	121/1	121/64	1102.64
7.1699	9	4:3	16/9	16/9	996.09
7.6147	15	7:2	49/4	49/32	737.65
7.8138	13	5:3	25/9	25/18	568.72
8.6439	13	5:4	25/16	25/16	772.63
8.7846	17	7:3	49/9	49/36	533.74
9.6147	17	7:4	49/16	49/32	737.65
9.8138	15	6:5	36/25	36/25	631.28
10.2586	21	7:5	49/25	49/25	1165.02
10.6439	15	8:5	64/25	32/25	427.37
11.6147	19	8:7	64/49	64/49	462.35
12.2586	23	10:7	100/49	50/49	34.98
12.9837	19	10:9	100/81	100/81	364.81
14.0888	29	12:11	144/121	144/121	301.27

Note that the two indices sometimes disagree in their fine ordering — 5:1 has lower Tenney distance than 3:2 but a higher Euler gradus (9 versus 7), because Euler penalises the prime 5 more heavily than Tenney does. **Where two defensible indices disagree, publish both and do not adjudicate.** Neither is a measurement.

The most consonant nontrivial essence pairs are exactly the octave-equivalent ones. Every pair whose index ratio reduces to `2:1` has `HD_ess = 2.0000`, gradus 3, and an interval of exactly two octaves — pitch class identical. Nothing else comes close: the next entry, `3:1`, is at 3.1699.

Which produces the central design tension of this chapter, and it is exact:

CONSONANCE IS REDUNDANCY. The pairs that sound most related are precisely the pairs that occupy the same pitch class — that is, the pairs the lattice cannot distinguish. **In the essence lattice, harmonic agreement and informational redundancy are the same phenomenon.**

This is the sonic image of Chapter 11's Parallel Wall, and it is not a metaphor: the arithmetic is the same arithmetic. A maximally consonant essence set is a maximally collapsed one. Searching all 6-element subsets of indices 1..12 exhaustively (924 subsets), by mean pairwise essence harmonic distance:

```
minimum mean HD_ess overall       4.0353   {1, 2, 3, 4, 6, 12}   2 distinct classes
maximum mean HD_ess overall      11.4064   {7, 8, 9, 10, 11, 12} 6 distinct classes
 minimum mean HD_ess with 6 classes   8.4731    {1, 3, 5, 7, 9, 11}
```

Read the first line. **The most consonant six-element essence set available collapses to two pitch classes.** Six essences, two distinguishable tones: u has gone to zero in the sonic register and the set is worth `ln 2` where it should be worth up to `ln 6`.

Read the third line. **Subject to keeping all six pitch classes distinct, the most consonant six-element set is the odd indices `{1, 3, 5, 7, 9, 11}`** — which is exactly the canonical class-representative set from Result Two. That is a satisfying convergence and it is also a usable design rule: **when a subset of a large essence set must be chosen for sonification, choose the odd indices.** They are the class representatives, they are guaranteed distinct, and among all fully-distinct choices they are the least dissonant.

FVP . Pitch-Class Folding by Odd Part

`class(n) = odd part of n`; number of classes over `1..N` is `ceil(N/2)`

Gate	Test	Result
FVP-1 DIMENSIONS	Both index and odd part are dimensionless integers; the class count is a dimensionless count. Consistent.	**PASS**

Gate	Test	Result
FVP-2 SCALE TYPE	The odd part is a NOMINAL label — a class identifier, not a magnitude. **Do not do arithmetic on class numbers.** The class's octave position (in cents) is interval-scaled and may be subtracted. The count `ceil(N/2)` is a ratio-scaled count. Three different scale types in one construction, and they must not be mixed.	**PASS with an explicit three-way scale warning**
FVP-3 LIMITING CASES	`N = 1`: one class, the ground. `N = 2`: still one class — indices 1 and 2 are two octaves apart and harmonically identical. `N -> infinity`: classes grow as `N/2`, so the vocabulary grows linearly while the pitch span grows only logarithmically, which is why collisions appear (§33.8).	**PASS — and the limiting behaviour explains the collisions**
FVP-4 FIXED POINTS	The odd indices are the fixed points of the folding map: `class(m) = m` iff `m` is odd. **The class representatives are exactly the odd indices**, which is the result §33.11 recovers independently by minimising dissonance.	**PASS — fixed point set is interpretable and useful**
FVP-5 IDENTIFIABILITY	The class does not identify the index: `{3, 6, 12, 24}` all map to class 3. Recovering `n` from its class requires the octave count `a`, i.e. the pair `(a, m)`, which is a bijection with `n`. **Publish `(a, m)`, not `m` alone, if the index must be recoverable.**	**PASS given (a, m); FAILS given m alone**

Gate	Test	Result
FVP-6 FALSIFICATION	A theorem of arithmetic (§33.3), proved and verified by enumeration for `N = 1..24`, all 24 agreeing. Not empirically falsifiable; arithmetically checkable and checked.	**N/A (theorem) — verification reported**
FVP-7 REPRODUCIBILITY	Exact and machine-independent. As with the essence law, reproducibility of the *index* transfers to the ordering (R3).	**PASS for the arithmetic; conditional on R3 for the input**

33.12 Constructing a defensible essence set

The requirements R1-R3 of §33.4 are the gate. This section is the procedure for passing it.

Step 1 — State the domain and its boundary. What is being categorised, and what is excluded. Without a boundary, exhaustiveness (R1) cannot be checked.

Step 2 — Identify the ordering source. Three kinds are admissible, in descending order of strength.

(a) An intrinsic physical or structural magnitude. The categories are ordered by a measurable quantity that exists independently of the analyst — distance, duration, count, energy, age. This is the strongest source because R2 is satisfied automatically: the ordering predates the analysis. The published proxemic zones used in the worked example are of this kind, ordered by distance.

(b) A generative or constructive order. The categories are produced by a process with a natural sequence, and the index is position in that sequence — stages of a protocol, steps of a manufacturing chain, layers of a stack. Admissible when the sequence is fixed by the process rather than by preference.

(c) A measured empirical rank with demonstrated reliability. The ordering is established by measurement — a Rasch-calibrated item difficulty ordering, for instance (Chapter 54) — with inter-rater or test-retest reliability reported. Weakest of the three, because the ordering is now an estimate carrying its own uncertainty, which must be propagated. **If the rank order is unstable across samples, the frequencies are unstable, and the whole construction must be reported with that instability visible.**

Step 3 — Fix the direction. Ascending or descending must be justified, because reversing the ordering completely changes every frequency and every interval. The default rule adopted here: **index 1 is the ground, so index 1 goes to the category that is nearest, smallest, earliest, or most basic in the ordering source.** Reversal is permitted only with a stated reason.

Step 4 — Count. `N` is the number of categories. It determines the ceiling `ln N` (Chapter 32) and the class count `ceil(N/2)` (§33.3). Both are reported.

Step 5 — Check for collisions. Compute the class cents and check the gaps (§33.8). Any gap under about 30 cents will be heard as beating rather than as an interval, and must be disclosed in the publication of the set.

Step 6 — Publish the whole construction, not the frequencies. The domain, the boundary, the ordering source, the direction, the index assignment for every category, `N`, the class count, the collisions, the band, and the transposition `j`. **A frequency published without its construction is not a result; it is a decoration.**

What is inadmissible as an ordering source: alphabetical order; order of discovery or historical accident; order of importance as judged by the analyst; frequency of occurrence in a convenience sample; and — the one that matters most — **any order arrived at by trying orderings and keeping the one that sounded best.** That last is the essence-set form of p-hacking. It converts the whole apparatus into an elaborate way of writing down a preference.

33.13 Band placement and transposition

Raw essence frequencies for a 24-set span 0.483 Hz to 278 Hz — from far below hearing to the low bass. Octave transposition brings them into a chosen band:

```
f_audible = f * 2^j          j an integer
```

Transposition by a whole number of octaves preserves pitch class exactly, which is the whole reason it is the house method. Placing indices 1..24 in the band `[110, 220)` Hz, one `j` per index:

n	f raw Hz	j	f*2^j Hz	class m	n	f raw Hz	j	f*2^j Hz	class m
1	0.483426	8	123.7570	1	13	81.698949	1	163.3979	13
2	1.933703	6	123.7570	1	14	94.751444	1	189.5029	7
3	4.350832	5	139.2266	3	15	108.770790	1	217.5416	15
4	7.734812	4	123.7570	1	16	123.756988	0	123.7570	1
5	12.085643	4	193.3703	5	17	139.710037	0	139.7100	17
6	17.403326	3	139.2266	3	18	156.629937	0	156.6299	9
7	23.687861	3	189.5029	7	19	174.516690	0	174.5167	19
8	30.939247	2	123.7570	1	20	193.370293	0	193.3703	5
9	39.157484	2	156.6299	9	21	213.190748	0	213.1907	21
10	48.342573	2	193.3703	5	22	233.978055	−1	116.9890	11
11	58.494514	1	116.9890	11	23	255.732213	−1	127.8661	23
12	69.613306	1	139.2266	3	24	278.453222	−1	139.2266	3

Read the class column against the frequency column: every index in class 1 lands on 123.7570 Hz, every index in class 3 on 139.2266 Hz. **Per-index transposition into a single octave collapses the set onto its pitch classes exactly — twelve tones for twenty-four essences, as the theorem requires.**

And that is the hazard. Per-index transposition destroys the intervals between indices while preserving classes. The 24-essence set spans 11003.91 cents — more than nine octaves — and folding it into one octave discards all of that. **Two different operations are available and they must never be confused:**

```
SHARED j   f_i * 2^j for one fixed j    preserves ALL intervals; needs a wide band
PER-INDEX j f_i * 2^{j_i}            preserves pitch CLASSES only; fits one octave
```

State the band and the `j` for every value published (CANON_SONIC §5), and state which of the two operations was performed. A set of values whose intervals are the point takes a shared `j`.

And a shared `j` has a hard capacity, which Chapter 71 computes exactly. The set `n = 1..N` spans a frequency ratio of `N^2`, so fitting it inside 20 Hz - 20 kHz under one shared `j` requires `k * 2^j >= 20` at the bottom and `k * N^2 * 2^j <= 20000` at the top. The first inequality forces `j >= 6`, since `20/k = 41.37` and `2^5 = 32` is too small. The second then caps `N`:

```
j = 6    k * 2^6 =    30.94 Hz     N <= 25      k * 25^2 * 2^6 = 19337.0 Hz
j = 7    k * 2^7 =    61.88 Hz     N <= 17
j = 8    k * 2^8 =   123.76 Hz     N <= 12
j = 9    k * 2^9 =   247.51 Hz     N <=  8
```

At most 25 essences can be rendered under a single shared transposition, and only at `j = 6`. That is a bound on set size, not on volume or equipment, and it does not move. It has a direct consequence for the sets this manual actually builds: **a twelve-house set fits comfortably at `j = 6` or `j = 8`; a thirty-two-system set does not fit under any shared `j` and must be rendered under split transposition** — two or more shared-`j` groups, each internally interval-preserving, with the relation *between* the groups explicitly declared and not silently implied. Splitting is a legal move; leaving the split undeclared is the failure, because a reader cannot tell a split rendering from a broken one without the declaration.

One further constraint applies to any rendering and belongs beside the band. A rectangular on/off gate has **no finite spectral second moment** — its `sinc^2` tails fall as `1/f^2` and the second-moment integral is numerically divergent — so it has no defined bandwidth and **is disqualified as a rendering envelope** (Chapter 71). Use a windowed envelope; Chapter 31 §31.9 gives the house default. A hard-gated tone is not the tone whose frequency you published.

FVP . Octave Transposition

`f_audible = f . 2^j`, `j` an integer

Gate	Test	Result
FVP-1 DIMENSIONS	Hz times a dimensionless power of two gives Hz. Consistent.	**PASS**
FVP-2 SCALE TYPE	`f` is ratio-scaled; multiplication by a positive constant is permissible and preserves all ratios. `j` is an integer count of octaves, ratio-scaled.	**PASS**

Gate	Test	Result
FVP-3 LIMITING CASES	`j = 0` is the identity. `j –> +infinity` leaves the audible band above; `j –> –infinity` below. Both are outside the working range, which is why the band must be declared.	**PASS with declared band**
FVP-4 FIXED POINTS	The pitch class is invariant under the operation for every `j`. **That invariance is the entire justification for the method** — and it holds only for integer `j`.	**PASS**
FVP-5 IDENTIFIABILITY	`f` is recoverable from `f_audible` only if `j` is published. Without `j` the raw essence frequency is unrecoverable and the value cannot be checked. **This is why CK-6 exists.**	**PASS if j published; FAIL otherwise**
FVP-6 FALSIFICATION	An exact operation, not an empirical claim. Checkable and checked in the table above.	**N/A (definition)**
FVP-7 REPRODUCIBILITY	Exact given `f`, band, and rule for choosing `j`. **The rule must be published, not just the `j`** — "lowest `j` placing `f` in `[110, 220)`" is a rule; "j = 4" is a value.	**PASS if the rule is published**

33.14 What a well-formed essence set looks like on paper

Summarising the engineering into a shape a reader can recognise:

- **Small.** Under about twelve. Above that, near-unison collisions appear that the ordering did not intend (§33.8), and the marginal ceiling gain is logarithmic anyway (Chapter 32).
- **Ordered from outside.** The ordinal comes from the domain, established before the sonification was contemplated (R2).
- **Fully enumerated in print.** Every category, its index, its class, its cents.
- **Honest about its collisions.** Any class gap under 30 cents disclosed.
- **Reported with `N`, `ceil(N/2)`, the ceiling `ln N`, the band, and the transposition rule.**
- **Free of adjustment.** No index was moved to improve the sound.

- **Badged throughout.** Essence frequency, not physical resonance; chosen `k`; postulated square law; ordinal input promoted by postulate.

The Specification

ES-1. An essence set is published with its **domain and boundary** stated.

ES-2. Categories are **exhaustive and disjoint** over that domain (R1), and the check is reported.

ES-3. The **ordering source** is stated and classified as (a) intrinsic magnitude, (b) generative sequence, or (c) measured rank with reliability. Type (c) reports its reliability statistic.

ES-4. The ordering is **established independently of the sonification** (R2). Any set whose ordering was chosen or adjusted for its sound is disqualified and may not be published as an essence set.

ES-5. The **direction** of the index (ascending from ground) is stated and justified.

ES-6. Assignment is reproducible (R3): the rule is stated in enough detail that a second analyst reproduces the index, and inter-rater agreement is reported where the assignment involved judgement.

ES-7. `N`, the class count `ceil(N/2)`, and the ceiling `ln N` are published together.

ES-8. The **full index table** is published: category, index `n`, factorisation `2^a.m`, class `m`, raw `f = k n^2`, class cents.

ES-9. Class collisions under 30 cents are disclosed, with the pair identified and the gap in cents.

ES-10. The **band and the transposition rule** are published, and whether a shared or per-index `j` was used (CANON_SONIC §5).

ES-11. Every published frequency carries: *this is an essence frequency, not a physical resonance.*

ES-12. `k` is published as a **chosen** reference constant, with the note that it cancels from every interval.

ES-13. The **scale promotion** (§33.5) is badged wherever intervals are given a quantitative interpretation: the ordering is data, the spacing is the model.

ES-14. Consonance indices are published with their definitions and their CHOSEN status. Where two defensible indices disagree in ranking, both are published and neither is adjudicated.

ES-15. No essence set is used to rank, score, select, or exclude a person. The index is a position in a category ordering, not a measure of worth.

ES-16. Anything delivered as tone is **audible, disclosed and refusable**, with no subliminal content of any kind (CANON_SONIC §7).

ES-17. Where a subset of a large set must be sonified, the **odd indices** are the default choice, being the class representatives and, among fully-distinct choices, the least dissonant (§33.11). A different subset requires a stated reason.

ES-18. Every formula in the construction carries its full seven-gate FVP table in print, including gates it fails or cannot yet attempt.

The Failure Modes

F-33.1 . THE TUNED ORDERING. *Signature:* an essence set whose index assignment was revised after the first listening. *Damage:* R2 is void; the frequencies encode the analyst's preference and nothing else. Every downstream interval is an artefact. *Detection:* ask for the date and provenance of the ordering relative to the first sonification. *Fix:* none available after the fact. Re-derive the ordering from an external source and discard the tuned version. **This is the essence-set form of p-hacking and it is disqualifying.**

F-33.2 . THE ARBITRARY ORDINAL. *Signature:* categories indexed alphabetically, by order of discovery, or by the order they appeared in a spreadsheet. *Damage:* the lattice is arbitrary; a different arbitrary ordering gives different frequencies with equal claim to correctness. *Detection:* apply a second plausible ordering and compare. If the intervals change materially, the ordering was doing no work. *Fix:* find an intrinsic magnitude or a generative sequence, or decline to sonify.

F-33.3 . THE OVERLAPPING CATEGORY. *Signature:* an entity in the domain that belongs to two categories. *Damage:* R1 fails; the entity has two indices, two frequencies, and a spurious interval between itself and itself. *Detection:* attempt to classify twenty randomly drawn cases. Any double-assignment is a hit. *Fix:* split or merge categories until the partition is clean, then recompute N and every index.

F-33.4 . THE CLASS COLLISION READ AS MEANING. *Signature:* two categories whose essence tones nearly coincide, interpreted as evidence that the categories are related. *Damage:* classes 3 and 17 are 6.00 cents apart because `9/8` and `289/256` are nearly equal. That is a fact about integers, not about the domain. **Reading it as a domain finding is pure numerology.** *Detection:* check whether the collision is predicted by §33.8's table for any set of that size. If it is, it is structural. *Fix:* disclose all structural collisions in advance (ES-9), so that none of them can later be mistaken for a result.

F-33.5 . THE OVERGROWN SET. *Signature:* an essence set with thirty or forty categories. *Damage:* the top of the lattice is congested (adjacent intervals under 150 cents), collisions multiply, the class count grows linearly while the audible span does not, and the ceiling gain from the extra categories is logarithmic and negligible. *Detection:* compute the adjacent-interval table. Any adjacent pair under about 100 cents indicates the set has outgrown the lattice. *Fix:* aggregate categories to a coarser, defensible partition and rebuild the index from the same ordering source.

F-33.6 . THE CONSONANCE CHASE. *Signature:* a set curated toward consonant intervals, or a subset chosen because it "sounds harmonious". *Damage:* consonance is redundancy in this lattice (§33.11). The most consonant six-set collapses to two pitch classes. **A beautiful-sounding essence set is very often a collapsed one, worth `ln 2` where it should be worth `ln 6`.** *Detection:* count distinct classes in the chosen subset. Fewer than `ceil(size/2)` means collapse. *Fix:* ES-17. Take the odd indices, accept the dissonance, and understand that the dissonance is the sound of the set actually carrying information.

F-33.7 . THE SCALE PROMOTION, UNBADGED. *Signature:* intervals from an ordinal-indexed set interpreted as measured magnitudes of relation. *Damage:* an unearned upgrade from rank

data to ratio-scale claim; the formula looks like a measurement and is a model. *Detection:* ask what would change if the spacing postulate were `n^1.5` instead of `n^2`. If the answer is "every interval" and the text does not say so, the badge is missing. *Fix:* ES-13, and the sensitivity note of Chapter 32 §32.6.

F-33.8 . THE LOST TRANSPOSITION. *Signature:* a published audible frequency with no `j` and no band. *Damage:* the raw essence frequency cannot be recovered, so the value cannot be checked or compared with another set. *Detection:* attempt to recompute `n` from the published Hz. If `sqrt(f/k)` is not near an integer for any small power of two, `j` is missing. *Fix:* ES-10. Publish the band, the rule, and whether `j` was shared or per-index.

F-33.9 . THE MIXED TRANSPOSITION. *Signature:* a set of values folded into one octave with per-index `j`, then compared by interval. *Damage:* all inter-index intervals are destroyed; a 1.15-octave relationship reads as 0.15 octaves. Comparisons across the set are meaningless. *Detection:* check whether `j` varies across the compared values. *Fix:* recompute with a shared `j` and a band wide enough to hold the span.

Worked Example

Building an essence set from a published, externally-ordered category system.

The domain and the ordering source

Hall (*The Hidden Dimension*, 1966) divides interpersonal distance into four zones — intimate, personal, social, public — and subdivides each into a close and a far phase, giving **eight categories ordered by physical distance.** Hall's nominal figures: intimate close 0 to 6 inches, intimate far 6 to 18 inches; personal close 1.5 to 2.5 feet, personal far 2.5 to 4 feet; social close 4 to 7 feet, social far 7 to 12 feet; public close 12 to 25 feet, public far beyond 25 feet.

Checking the requirements:

- **R1 (exhaustive, disjoint):** the zones tile the distance axis from zero upward without gaps or overlaps. **PASS**, subject to the boundary being "interpersonal distance in a face-to-face encounter", which is stated.
- **R2 (ordering from outside):** the ordering is by **measured distance**, an intrinsic physical magnitude, published in 1966 with no relation to any sonification. This is ordering source type (a), the strongest. **PASS.**
- **R3 (reproducible assignment):** given a measured separation, the zone assignment is mechanical. **PASS.** Note that Hall's zone *boundaries* are nominal and culturally variable — Hall said so — but the *ordering* is not in question, and only the ordering is used here.

Direction: index 1 is the ground, assigned to the nearest zone, per the default rule of §33.12 Step 3.

The set

`N = 8`. Class count `ceil(8/2) = 4`. Ceiling `ln 8 = 2.079442` nats = **3.000000 bits exactly** = 6 octaves exactly (Chapter 32).

n	category	f = k n^2 Hz	class m	class ratio	class cents	j	f∗2^j Hz
1	intimate close	0.483426	1	1/1	0.000	8	123.7570
2	intimate far	1.933703	1	1/1	0.000	6	123.7570
3	personal close	4.350832	3	9/8	203.910	5	139.2266
4	personal far	7.734812	1	1/1	0.000	4	123.7570
5	social close	12.085643	5	25/16	772.627	4	193.3703
6	social far	17.403326	3	9/8	203.910	3	139.2266
7	public close	23.687861	7	49/32	737.652	3	189.5029
8	public far	30.939247	1	1/1	0.000	2	123.7570

Band [110, 220) Hz, per-index j, rule: the unique integer j placing f . 2^j in the band. **Every value in the last column is an essence frequency, not a physical resonance. No claim is made that any distance, person, or encounter vibrates at any of these values.**

Reading the result

Four of the eight zones share one pitch class. Indices 1, 2, 4 and 8 all have odd part 1 — intimate close, intimate far, personal far, and public far are octave-identical in this lattice, all landing on 123.7570 Hz in the chosen band. The distinct classes are {1, 3, 5, 7}, exactly the four predicted by the theorem, and exactly the odd indices predicted as class representatives by §33.11's fixed-point result.

This is a genuine structural finding about the lattice and it is emphatically not a finding about proxemics. Nothing about intimate distance and public distance is being claimed to be similar. The collapse is caused by the indices being powers of two, which is caused by where those four zones happened to fall in Hall's ordering. **Had Hall subdivided into three phases per zone rather than two, giving twelve categories, the class structure would be entirely different — and Hall's choice was about ethnography, not about arithmetic.**

That sentence is the whole lesson of the worked example. **The lattice's structure is a fact about the index, and the index's meaning is a fact about the domain, and the two must never be run together.**

Collision check (ES-9)

The four realised classes sit at 0.000, 203.910, 737.652 and 772.627 cents. Gaps:

```
    class 1 (   0.000)  ->  class 3 ( 203.910)      203.910 cents
    class 3 ( 203.910)  ->  class 7 ( 737.652)      533.742 cents
     class 7 ( 737.652)  ->  class 5 ( 772.627)        34.976 cents  <--
COLLISION, disclosed
    class 5 ( 772.627)  ->  class 1 (1200.000)      427.373 cents
```

Classes 7 and 5 — public close and social close — are 34.976 cents apart and will beat rather than form an interval if sounded together. Disclosed in advance, as ES-9 requires, precisely so that no one can later report it as a finding about the relationship between social and public distance. It is 49/32 against 25/16, and it is arithmetic.

The pairwise interval table

pair	interval (cents)	essence ratio	HD_ess	pair	interval	ratio	HD_ess
1–2	2400.00	4/1	2.0000	3–5	1768.72	25/9	7.8138
1–3	3803.91	9/1	3.1699	3–6	2400.00	36/9	2.0000
1–4	4800.00	16/1	4.0000	3–7	2933.74	49/9	8.7846
1–5	5572.63	25/1	4.6439	3–8	3396.09	64/9	9.1699
1–6	6203.91	36/1	5.1699	4–5	772.63	25/16	8.6439
1–7	6737.65	49/1	5.6147	4–6	1403.91	36/16	5.1699
1–8	7200.00	64/1	6.0000	4–7	1937.65	49/16	9.6147
2–3	1403.91	9/4	5.1699	4–8	2400.00	64/16	2.0000
2–4	2400.00	16/4	2.0000	5–6	631.28	36/25	9.8138
2–5	3172.63	25/4	6.6439	5–7	1165.02	49/25	10.2586
2–6	3803.91	36/4	3.1699	5–8	1627.37	64/25	10.6439
2–7	4337.65	49/4	7.6147	6–7	533.74	49/36	10.7846
2–8	4800.00	64/4	4.0000	6–8	996.09	64/36	7.1699
3–4	996.09	16/9	7.1699	7–8	462.35	64/49	11.6147

All 28 pairs. The four most consonant nontrivial pairs are 1-2, 2-4, 3-6 and 4-8, all at `HD_ess = 2.0000` — and all four are index ratios of `2:1`, hence pitch-class collisions. **The most consonant pairs are exactly the redundant pairs, as §33.11 requires. The lattice is behaving.**

A negative example, for contrast

Consider building an essence set from this manual's own six storage-law parameters — `delta, eta, N, u, C, Q`. It fails at Step 2 and it fails hard.

- **R1:** the six are not categories of a common domain at all. `N` is a count, `u` is a geometric quantity in `[0,1]`, `Q` is a stock, `C` is a flow. They do not partition anything. **FAIL.**
- **R2:** no intrinsic magnitude orders them. They could be ordered by the sequence in which the storage law is written, by alphabet, or by importance — three different orderings with no external arbiter, each giving a completely different lattice. **FAIL.**
- **R3:** two analysts would not agree. **FAIL.**

Therefore the six parameters get no frequencies. They can be tabulated, reasoned about, measured, and used — everything this manual actually does with them — but they cannot be sonified, because there is no defensible ordinal and the law would convert an arbitrary choice into a confident-looking set of Hertz values.

That is the point of the requirements, and it is why they were stated before the arithmetic rather than after. Most category sets do not qualify. The instrument is narrow, and a narrow instrument used honestly is worth more than a wide one used loosely.

Chapters 32 and 33 have supplied the units and the lattice. Part VII turns to the register these units were built for: the perceptual lens itself, and the moment it produces.

Part VII — The Perceptual Lens

Chapter 34 . The Moment Equation

Volume One built the outside. Rooms, tables, cadences, institutions — the circumstance in which stored order accumulates. Every chapter of it ran on one recursion, `Q_{t+1} = (1−δ)Q_t + ηS_t − C_t`, and every chapter of it took the input `S_t` as something a designer arranges by moving people and furniture around.

Volume Two asks the question that arrangement cannot answer. **A room is not the same room to two people standing in it, and the storage law does not integrate the room — it integrates whatever arrived.** Something stands between the environment and the input. This chapter names it, gives it an equation, and submits that equation to the full seven-gate validation protocol.

The equation is the author's own, and it is the centre of this volume:

```
M = (D / P) . E
```

Perspective over perception, times environment, equals moment. Four symbols. It is small enough to write on a card and it takes the rest of this chapter to say honestly.

The Law

34.1 The statement, badged

> **THE MOMENT EQUATION [POSTULATE — the author's formula].** For an observer in contact with an environment, the moment as lived is the environment scaled by the ratio of the frame the observer brings to the intake the observer takes:
>
> ```
> M = (D / P) . E
>
> D = perspective : the structure the observer imposes — model, expectation,
>                   prior, doctrine, training, the shape the observer is
>                   prepared to find
> P = perception  : the structure the observer takes in — intake, evidence
>                   actually registered, the raw admission of what is there
> E = environment : what is actually available to be had at this place and time
> M = moment      : the situation as it is lived; the input the observer's own
>                   storage law will integrate
> ```
>
> This is a postulate, not a theorem. It is not derived from Papers A–E. It is a proposed structure for a quantity nobody has yet measured directly, and it earns its place in this manual only by passing FVP-1 through FVP-5 in print and by naming, in §34.24, exactly what would kill it.

Nothing in Volume One depends on this equation. Nothing proved in Papers A–E is affected by whether it survives. **It is offered as the load-bearing hypothesis of the internal instrument, and it is offered with the badge on, because CANON §5 forbids presenting a postulate without one.**

34.2 Why a ratio and not a difference

The first structural decision in the formula is that D and P are combined by division rather than by subtraction. That decision is not stylistic and it is not free. It is the decision that puts the Moment Equation into the same register as everything else in this volume.

Four quantities in this framework are logarithmic (CANON_INTERNAL §1): surplus `F = ln N_eff`, pitch `p = log2 f`, sensation `S = k ln(I/I_0)`, and the moment. Taking logs of the Moment Equation:

```
ln M = ln D − ln P + ln E
```

A ratio in the world is a difference in the lens. That single line is the whole reason the formula divides. If D and P were combined additively, the moment would not compose with surplus, with pitch, or with sensation, and none of the conversion keys of Chapter 32 would apply to it. By dividing, the moment joins the register the other three already occupy, and every operation available there — the log-ratio distortion measure of Chapter 35, the geometric mean, the log-space contraction — becomes available here.

There is a second reason, and it is empirical rather than architectural. Weber's law says the just-noticeable difference is a constant *ratio*, not a constant difference. If the instrument being modelled reports ratios, the model should be built out of ratios. Chapter 36 makes this argument in full and shows that it forces the form.

34.3 FVP-1 — Dimensions

```
[M] = ([D] / [P]) . [E]
```

For the expression to be well formed, `[D] = [P]`. Their ratio is then dimensionless, and

```
[M] = [E]
```

The moment inherits the units of the environment. This is not a convenience; it is the formula's central assertion. M and E are the same kind of thing measured on the same scale, and the observer's contribution is a pure dimensionless multiplier on it. A moment is not a different species of object from the environment that produced it. It is that environment, scaled.

This immediately constrains what D and P are allowed to be. They must be commensurable — the same quantity, measured the same way, once for the frame and once for the intake. **You may not put "years of training" over "number of things noticed" and call the result a ratio.** The operational requirement that falls out of this is stated as S34.1 and demonstrated in the worked example: D and P must both be *recorded items per unit of environment present*, or both be *precision in the same units*, or both be *bits of asserted structure per bit of available structure*. Anything else fails FVP-1 and must not be written as a fraction.

For the units of E and therefore of M, this manual adopts **nats of usable structure**, so that M plugs directly into the storage law as `S_t` and inherits the ceiling `0 <= M <= ln N` derived in §34.19. Other choices are possible and must be declared.

34.4 FVP-2 — Scale type

This is the gate the formula is most likely to fail in careless hands, and Chapter 53 is devoted to it. Stated here in short:

Operation	Minimum scale required
Ranking, ordering	Ordinal
Addition, subtraction, mean	Interval
Multiplication, division, ratio	**Ratio, with a true zero**

`M = (D/P).E` contains one division and one multiplication. **Therefore D, P and E must each be measured on a ratio scale with a meaningful zero, or the equation is not an equation.** If the best available data on D is a five-point self-report of "how strongly did you expect this", the ratio D/P is meaningless: a 4 is not twice a 2 on a Likert item, because a Likert item has no true zero and no defensible unit.

The honest consequence, and this manual takes it:

> **DEMOTION RULE.** Where only ordinal data exist for D or P, the Moment Equation shall be written and read as a *heuristic ordering* — "more frame relative to intake produces a larger and more model-shaped moment" — and shall not be printed as a fraction, computed, or used to generate a number. The demoted form is still useful. It is not the equation.

The route out of demotion is not rhetorical, it is procedural, and psychophysics has already walked it: Stevens' method of magnitude estimation produces genuine ratio-scale measurements of subjective quantity from human subjects, and it replicates (Chapter 38). The worked example in §34.25 takes a different and cheaper route — it counts things — but the count route has its own limits, stated there.

And the rule above is not hypothetical. It fires now, on this chapter's own formula, and the manual applies it to itself before it applies it to anyone else. Chapters 52 and 53, working the protocol through in full, return **DEMOTE on FVP-2 for `M = (D/P).E` as the equation currently stands.** No instrument yet exists that puts D and P on a ratio scale with a shared true zero. Therefore, and this is stated at the first presentation of magnitude rather than in a footnote at the end:

> **WHAT IS NOT REPORTABLE AS A MEASURED QUANTITY.** The *magnitude* of `D/P`; the distortion `Delta = |ln r|` as a number; and the quantitative form of the join `M_t -> Q_{t+1}` (Chapter 40). Every figure of this kind printed in this chapter and the next is an **illustrative construction demonstrating the structure of the formula**, and is labelled as such.
>
> **WHAT SURVIVES INTACT.** The four limiting cases (§34.9–§34.12); the fixed point `D = P` and its uniqueness (§34.13); and the ordering *over-framed* / *calibrated* / *under-framed* (§34.5). **All three need only ordinal information, and none of them is touched by the demotion.** What is demoted is magnitude, not structure.

The constructive route out is specific, principled, and costed. Chapter 54, **Theorem 54.3**, given badged Postulate SO1: if D and P are measured as Rasch person-measures on one common scale, then `Delta = |ln(D/P)| = |theta_D − theta_P|` is genuinely ratio-scaled, because Rasch separation supplies the true zero that division requires and fixes the unit up to choice. The measurement cost is stated in §54.12 and it is the honest part: **approximately 123 items per term to establish a half-logit distortion at 95 per cent.** A twelve-item pair of scales

cannot resolve anything below about 1.6 logits and a practitioner reporting `r = 1.2` from one is reporting noise.

Read this as the protocol working, not as a caveat. A validation gate that demotes its author's own keystone formula, names the exact quantity it kills, prices the study that would revive it, and leaves the structural content standing is doing precisely the job it was written to do. The alternative — printing `r = 1.33` and hoping no one asks what the unit is — is the practice this volume exists to replace.

34.5 The composite parameter

The formula has four symbols but three degrees of freedom, because D and P enter only through their ratio. Define

```
r = D / P          the CALIBRATION RATIO      (dimensionless)

M = r . E
```

Everything that D and P do to the moment, they do through r. **A tenfold increase in both is a tenfold increase in nothing.** This is the single most important structural fact about the equation after the fixed point itself, and it has an immediate practical consequence: an intervention that raises frame and intake together changes what the observer can handle, not how distorted the observer is. Chapter 35 is entirely about r; this chapter establishes what r does inside the equation.

The signed distortion is `ln r`, the unsigned distortion is `|ln r|`, and both live in Chapter 35.

34.6 The full FVP table

`FVP TABLE — THE MOMENT EQUATION  M = (D/P) . E`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units of each term; both sides agree; ratios dimensionless.	Requires `[D]=[P]`; then `r` is dimensionless and `[M]=[E]`. Manual adopts nats of usable structure for E and M. Any D, P pair not measured in identical units is inadmissible.	**PASS**, conditional on the commensurability requirement S34.1

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Stevens' typology per term; division requires ratio scale with true zero.	D, P, E must all be ratio-scaled. **No instrument yet puts D and P on a common ratio scale, so the magnitude of `D/P`, the number `Delta = |ln r|`, and the quantitative form of the `M -> Q` join are not reportable as measured quantities (Ch 52, Ch 53).** The limiting cases, the `D = P` fixed point and the three-way ordering survive, needing only ordinal information. Route out: **Ch 54 Thm 54.3**, Rasch logit difference, given badged Postulate SO1 — at ~123 items per term for a half-logit distortion.	**DEMOTE** — structure stands, magnitude does not; **Rasch route specified and costed**
FVP-3 LIMITING CASES	Behavior as each term -> 0 and -> infinity; every divergence must name a real phenomenon.	Four cases, §34.9–§34.18. `P–>0` diverges and names dogma/apophenia/hallucination. `D–>0` gives 0 and names undifferentiated sensation. `E–>0` gives 0. `r–>infinity` and `r–>0` both name observed clinical states. No unexplained divergence.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	Solve for the invariant; state its meaning.	`D = P => r = 1 => M = E`, for every E. Undistorted contact. Under the correction map of §34.16 it is the unique fixed point and is attracting for gain `0 < a < 2`.	**PASS**
FVP-5 IDENTIFIABILITY	Can parameters be recovered from obtainable data? Do distinct parameter sets predict identical observations?	**Partially. `r` is recoverable; D, P and E are not separately recoverable from M alone.** `(D, P, E)` and `(kD, kP, E)` are observationally identical for any `k > 0` — a one-parameter unidentified family. The count-based estimator of §34.25 recovers `r` directly and cheaply. E requires an independent instrument.	**OPEN — r identified, the D/P split is not, E requires external measurement**
FVP-6 FALSIFICATION	Preregister; fit on training; freeze; report held-out; beat persistence, direct-input, matched autoregressive, and **shuffled-input** baselines. State the killing result.	Protocol specified in §34.24. Killing result stated: if a frozen `r` estimated at time t fails to beat the shuffled-input control at predicting time t+1's error signature, the formula has no content.	**NOT RUN — protocol stated, study not performed**

Gate	Test applied	Finding	Verdict
FVP-7 REPRODUCIBILITY	Same operationalization, different people, different sites: do the parameters recover?	The count-based estimator is procedure-portable in principle: the seeded-item audit requires no instrumentation beyond a defect list. **Whether `r` is a stable trait-like parameter across occasions and raters, or a property of the particular item set, is unknown and is the decisive question.**	**NOT RUN — and this is the gate that decides whether any of this is science**

One gate of seven is demoted and two are unrun, and this manual does not pretend otherwise. The equation is structurally sound, dimensionally consistent, has a meaningful fixed point, and has an estimator that can be computed from counts on a Tuesday afternoon. **What it does not yet have is a unit.** Gate 2 returns DEMOTE — magnitudes derived from `D/P` are not reportable as measured quantities until D and P are co-calibrated, and Theorem 54.3 says exactly how that is done and what it costs. The formula has not been tested against a shuffled-input control by anyone, and its cross-individual recovery is unknown. Chapters 52 through 56 exist because of these rows of the table, and the demotion in the second row is the clearest demonstration in the volume that the protocol is not decorative.

34.7 The honesty clauses for this chapter

1. **The Moment Equation is a postulate, not a theorem.** It is not in Papers A–E and nothing in Papers A–E depends on it.
2. **The correspondence with predictive processing (§34.20) is a correspondence between formalisms.** Predictive processing is an active and contested research program. Structural similarity between two formalisms is not identity, and is not evidence that either is true.
3. **M is the moment as lived, not the moment as true.** The formula does not compute what is really there. It computes what the observer ends up with. §34.22 makes the difference quantitative and it is the most uncomfortable section in the chapter.
4. **The ceiling `M <= ln N` (§34.19) is inherited from the surplus law's ceiling and from the normalization choice that maps M into the surplus register.** It is a consequence of a convention plus a theorem, not a free-standing discovery.
5. **No claim is made that D, P or E correspond to anything neural.** They are parameters of a descriptive model of behavior.

The Engineering

34.8 Reading the equation before analysing it

Before the limits, read the equation the way you would read a circuit diagram.

E is the source. It has whatever it has: a room with three exits, a manuscript with eleven errors, a patient with a set of signs, a market with a set of prices. E does not depend on the observer, and the manual is committed to that — **the environment term is what makes this a formula about contact rather than a formula about opinion.** If E were observer-dependent the equation would be circular and would predict nothing.

P is the aperture. It is how much of E is admitted. It cannot exceed what is there and it is usually well under it.

D is the gain applied to whatever came through. It is the structure the observer is prepared to assert — the pattern-library, the training, the doctrine, the diagnosis they came in expecting.

M is what comes out the other end and gets integrated by the storage law. It is the thing you actually remember, the thing that becomes character, the thing your institution writes down.

The four limiting cases are what happens when one of those four elements goes to an extreme, and each of them is a named condition that clinicians, editors, auditors and inspectors already recognise without an equation. **The formula's claim to seriousness is that it predicts all four from one line, and that the four are not four separate pathologies but one parameter at four values.**

34.9 Limit 1: P -> 0. The frame with no intake

```
P –> 0,  D and E fixed  =>  M –> infinity
```

Computed, with D = 1 and E = 1:

P (intake)	M
1.0	1.000
0.5	2.000
0.25	4.000
0.1	10.000
0.05	20.000
0.01	100.000
0.001	1 000.000
0.000001	1 000 000.000

This is the only divergence in the formula, and FVP-3 requires that a divergence correspond to a real, nameable phenomenon or the formula is wrong. It does, and the phenomenon has several names depending on which discipline noticed it:

- **Dogma.** The doctrine is complete; no observation is admitted that could revise it; every event is read as confirmation. The moment is enormous and entirely manufactured.
- **Apophenia.** Structure is perceived in noise. The intake carries no signal; the frame supplies all of it; the result is a vivid, detailed, confident reading of nothing.

- **Hallucination, in the technical sense of a percept without a corresponding stimulus.** The formula does not explain hallucination and makes no clinical claim. It says only that a percept with a frame and no intake is exactly the term structure `D/0`.
- **The over-fitted institution.** Chapter 48's subject. The organisation's model is so elaborate and its measurement so degraded that everything that happens is already in the plan.

Note the direction of the divergence carefully, because it is counterintuitive and it is the formula's sharpest prediction. **Reducing intake does not reduce the moment. It inflates it.** The person who has stopped looking does not experience less; they experience *more*, and more confidently, and with less friction. A shrinking P feels like clarity. This is why the condition is stable, why it is pleasant, and why it does not self-correct: the internal signal that would tell you to look harder — the sense that the situation is thin — is precisely the signal that goes away as P falls.

The subjective symptom of P -> 0 is not confusion. It is certainty.

There is an honest limit to what the formula can say here. A divergence to infinity is a modelling artifact: no real observer's moment is unbounded. §34.19 supplies the ceiling `M <= ln N` that a real observer is subject to, and with that ceiling in place the correct statement is that as `P -> 0` the moment saturates at the observer's full repertoire — the observer sees the maximum their categories permit, all of it supplied by themselves. **Saturation at maximum internal structure with zero external content is a better description of dogma than divergence to infinity, and it is what the bounded form of the model actually predicts.**

34.10 Limit 2: D -> 0. The intake with no frame

```
D -> 0,  P and E fixed  =>  M -> 0
```

D (frame)	M (with P = 1, E = 1)
1.0	1.000000
0.5	0.500000
0.1	0.100000
0.01	0.010000
0.001	0.001000
0	0.000000

The moment collapses. Everything is arriving; nothing is being made of it.

This limit contains the chapter's second thesis, and it is stated as flatly as it can be stated: sensation is not experience. An observer can be saturated with intake — every detail registered, nothing missed, complete admission of the environment — and have no moment at all, because a moment requires something to be made of what arrived, and D is the term that makes it.

The named conditions:

- **Flooding.** The novice on their second week: they see everything, they can rank nothing, the four hours produce twelve lines. Chapter 35's worked example measures exactly this.
- **Overwhelm.** Not too much input in an absolute sense — too much input relative to the structure available to organise it. The distinction matters enormously for intervention, because the two diagnoses recommend opposite treatments. Reducing the input load treats the first. Supplying a frame treats the second, and the second is far more common.

- **Depersonalisation and the flattened moment,** described phenomenologically as things being present but not mattering. The formula offers a structure for that report — full P, near-zero D, therefore near-zero M — and makes no clinical claim whatever. It is a structural description of a reported experience, offered as such.
- **The data-rich, decision-poor organisation.** Chapter 48. Everything logged, nothing concluded.

The asymmetry between Limit 1 and Limit 2 is worth staring at. **Losing your intake makes the world feel bigger; losing your frame makes it feel like nothing.** Neither of them feels like an error from inside. One feels like insight and the other feels like exhaustion.

34.11 Limit 3: E -> 0. The environment with nothing in it

```
E -> 0,  r fixed  =>  M -> 0
```

With r = 3 fixed:

E	M
10	30.00
1	3.00
0.1	0.30
0.01	0.03
0	0.00

This limit looks trivial and is not. It says that **no calibration ratio, however good, manufactures a moment out of an environment that has nothing in it.** The best-tuned observer in an empty room gets an empty moment. r = 1 multiplied by zero is zero.

This is the formula's answer to a doctrine this manual explicitly rejects: the doctrine that attitude is sufficient. It is not sufficient, it is a multiplier, and a multiplier on zero is zero. The whole of Volume One — the room, the table, the sightlines, the acoustics, the cadence — is the engineering of E, and Volume Two does not replace it. **Volume Two is a multiplier on Volume One, and that is the exact relationship the equation encodes.**

Note also the reverse reading, which is the more useful one operationally. If M is small and you have established that r is near 1, then the problem is E and no amount of internal work will fix it. **A calibrated observer in a poor environment is not a psychological case. They are a correct instrument reporting a true reading, and the intervention is in the room.** More people are in this condition than in either of the distortion conditions, and mislabelling it as a perception problem is one of the most common and most damaging errors this manual can help you avoid.

34.12 Limit 4: the ratio limits

The three limits above hold two terms fixed and move one. The fourth is the composite:

```
r -> infinity  =>  M / E -> infinity     over-framed without bound
r -> 0         =>  M / E -> 0            under-framed without bound
r = 1          =>  M = E                 the fixed point
```

The full correspondence:

r	M/E	signed distortion ln r	Condition
0.01	0.01	-4.6052	Extreme under-framing
0.10	0.10	-2.3026	Extreme under-framing
0.25	0.25	-1.3863	Severe under-framing
0.50	0.50	-0.6931	Severe under-framing
0.80	0.80	-0.2231	Mild under-framing
0.90	0.90	-0.1054	Mild under-framing
1.00	**1.00**	**0.0000**	**Calibrated — contact**
1.10	1.10	+0.0953	Calibrated (band edge)
1.25	1.25	+0.2231	Mild over-framing
1.50	1.50	+0.4055	Marked over-framing
2.00	2.00	+0.6931	Severe over-framing
4.00	4.00	+1.3863	Severe over-framing
10.00	10.00	+2.3026	Extreme over-framing
100.00	100.00	+4.6052	Extreme over-framing

Read the third column, not the first. **The log-ratio is symmetric — r = 2 and r = 0.5 are the same distance from contact, 0.6931 nats, in opposite directions — and the raw ratio is not.** On the raw scale, over-framing appears to have unlimited room and under-framing appears squeezed into the interval (0,1). That is an artifact of the coordinate, and it is exactly the artifact Chapter 35 removes.

34.13 THE FIXED POINT: D = P

Everything above is preparation. This section is the content.

```
D = P    =>    r = 1    =>    M = E        for every E
```

When the frame the observer brings equals the intake the observer takes, the moment equals the environment. That is veridical perception, stated as an equation, and it is the most important line in Volume Two.

Take it apart slowly, because it is easy to read past.

It is an identity, not an approximation. At D = P the observer's contribution to the moment is exactly the multiplicative identity. Not "small". Not "negligible". Exactly one. The observer is still fully present — D is not zero, P is not zero, both are working at whatever magnitude the situation demands — and yet the product of their activity is the environment, unmodified.

It holds for every E. This is what makes it a *fixed point* of a relation rather than a coincidence at one value. There is no environment for which calibration fails to deliver contact, and no environment in which the calibrated observer must trade fidelity for scale. The condition is scale-free in E, which is the same structural property CANON §1.4 claims for the storage law and for the same reason: only the dimensionless part of the model carries the observer's contribution.

It is not the absence of a frame. This is the most common misreading and it inverts the formula. `D = 0` is not contact; `D = 0` is Limit 2, and it gives `M = 0`. **Undistorted contact is not the state**

of having no expectations. It is the state of bringing exactly as much structure as you took in. The tradition that says perception should be emptied of concept is describing `D -> 0`, and by this formula that tradition is describing not clarity but the abolition of the moment. What clarity requires is not less frame. It is *matched* frame.

It is not passivity, and it is not agreement. An observer at r = 1 in an environment containing a hostile fact perceives a hostile fact at full magnitude. Contact is not comfort. The fixed point makes no promise about what the moment contains; it promises only that the moment contains what the environment does.

It is not certainty. At r = 1 the observer asserts exactly what they took in, which in a half-legible environment is a half-legible reading. **Calibration means your confidence tracks your intake, which in a murky room means being appropriately unsure.** The person at r = 1 in fog reports fog. The person at r = 4 in fog reports a building.

34.14 Why D = P is a fixed point and not merely an equality

An equality is a coincidence at a value. A fixed point is the invariant of a map. The Moment Equation as written is a static relation, so to speak of a fixed point at all we need a dynamics, and the dynamics is the observer's own correction behavior.

Postulate the simplest correction the observer could run: after each contact, compare what you asserted with what came back, and move your frame toward your intake by a fraction of the discrepancy. In log coordinates, with correction gain a:

```
ln r_{t+1} = (1 − a) . ln r_t                THE CORRECTION MAP  [POSTULATE]

 fixed point:  ln r* = 0,  i.e. r* = 1,  i.e. D = P
```

That is a first-order linear recursion in log space, and it is the same object as the storage law of Chapter 2 with the input and cost set to zero: a leaky integrator relaxing to its fixed point. Its pole is `(1−a)`. By the contraction mapping theorem — **Banach, exactly as in CANON §1.3, and this is the same theorem doing the same work at a different scope** — it has a unique fixed point and iteration converges to it geometrically whenever `|1−a| < 1`, that is, whenever

```
0 < a < 2                        the stability condition for calibration
```

Computed trajectories from an initial distortion of `ln r = +0.8122`, the value measured in the worked example of §34.25:

gain a	pole |1-a|	steps to |ln r| < 0.10	first six values of ln r
0.10	0.90	20	+0.7310 +0.6579 +0.5921 +0.5329 +0.4796 +0.4316
0.20	0.80	10	+0.6498 +0.5198 +0.4158 +0.3327 +0.2661 +0.2129
0.35	0.65	5	+0.5279 +0.3432 +0.2231 +0.1450 +0.0942 +0.0613

gain a	pole |1-a|	steps to |ln r| < 0.10	first six values of ln r
0.50	0.50	4	+0.4061 +0.2031 +0.1015 +0.0508 +0.0254 +0.0127
0.75	0.25	2	+0.2031 +0.0508 +0.0127 +0.0032 +0.0008 +0.0002
1.00	0.00	1	0 0 0 0 0 0 (deadbeat)
1.30	0.30	2	-0.2437 +0.0731 -0.0219 +0.0066 -0.0020 +0.0006
1.80	0.80	10	-0.6498 +0.5198 -0.4158 +0.3327 -0.2661 +0.2129
2.00	1.00	never	-0.8122 +0.8122 -0.8122 +0.8122 -0.8122 +0.8122
2.30	1.30	never	-1.0559 +1.3726 -1.7844 +2.3197 -3.0157 +3.9204

Three readings from this table, all of which Chapter 35 develops into procedure:

1. **Under-correction (a small) converges, slowly.** The person who adjusts a little each time gets there. It takes twenty encounters at a = 0.10 and five at a = 0.35.
2. **Over-correction past `a = 1` oscillates** — the observer overshoots to the other side of contact each time — **but still converges as long as `a < 2`.** Oscillating correction is inefficient, not fatal.
3. **At `a >= 2` calibration diverges.** The observer who reverses their whole position at every contradiction, and then reverses it again, gets further from contact with every iteration. The trajectory at a = 2.30 goes 1.06, 1.37, 1.78, 2.32, 3.02, 3.92 — a person becoming more extreme in alternating directions, each swing driven by the last. **This is a mechanism, not a metaphor, and it is derived from the same stability condition an engineer applies to a feedback loop.**

The fixed point D = P is therefore not a static ideal but an attractor, and its basin is the set of correction gains strictly between 0 and 2. With no correction at all (a = 0) it is neutrally stable — whatever distortion you have, you keep. That is the default. Contact is maintained, not attained.

34.15 The fixed point in the correspondence table

The internal/external master table of CANON_INTERNAL §5 gives the D/P row two readings:

Term	INTERNAL (psychology)	EXTERNAL (sociology)
`D/P`	frame-to-intake ratio	**doctrine-to-data ratio**
`M`	the moment as lived	the situation as the group constructs it

The fixed point survives the change of scope unchanged, and that is the strongest single piece of evidence for the volume's thesis that psychology and sociology are one subject at two scopes. An institution at D = P asserts exactly what its instruments measured. It does not manufacture and it does not drown. Chapters 47 and 48 develop this at length; what belongs here is only the observation that `r = 1` reads correctly at both scopes without modification, which a merely metaphorical correspondence would not do.

34.16 What the fixed point costs

A fixed point that is free is a fixed point nobody needs a manual for. This one is not free, and the cost has a specific shape.

Maintaining `D = P` requires that every increase in intake be matched by an increase in frame, and every decrease likewise. Since both are effortful — P costs attention, D costs the construction and maintenance of usable structure — **calibration is a continuous matching problem under a load budget, and the load appears in the storage law as C.**

This produces the manual's least comfortable prediction and one worth stating in the open: **distortion is cheap and contact is expensive, in both directions.** Dropping P while holding D is the cheapest possible move — it costs less attention and produces a larger, more confident moment. That is the entire economics of dogma in one sentence. Dropping D while holding P is also cheap: you stop having to conclude anything. Both are downhill. Only the match is uphill.

The engineering response is the one the whole manual gives: if a state is uphill, **build the circumstance that holds it**, rather than requiring the person to hold it by effort. The specification in §34.26 is a list of such circumstances.

There is a second cost, and it is structural rather than effortful. It is stated here as a correction to the reading of §34.13 that this chapter would otherwise invite. At `r = 1` the input to the storage law is at its ceiling, `S = ln N`. Put that into Paper E Theorem 3.5 and then read the master gauge:

```
S = ln N                    (contact at the fixed point)
Q_inf = (eta ln N − C) / delta
h     = eta ln N − delta Q_inf − C
      = eta ln N − (eta ln N − C) − C
      = 0        exactly, by cancellation
```

A perfectly calibrated observer converges to exactly zero headroom. Not approximately, and not through mismanagement: perfect contact drives the input to its maximum, the maximum input drives the store to its maximum, and the maximum store consumes the whole budget in maintenance. This is the same identity Chapter 6 derives at external scope and Chapter 47 §47.6 derives for institutions, arriving here from the perceptual side.

> **CORRECTION, and it applies wherever this chapter or Chapter 35 speaks of calibration as a goal: `r = 1` is NOT an unqualified target.** Calibration maximises fidelity and veridical throughput (§34.22). It does not maximise room to build, and at equilibrium it leaves none. An observer who wants headroom at a settled state must either shed `Q` deliberately (Chapter 40 §40.9) or accept `r != 1` as a declared policy (Chapter 35 §35.20). **The manual states the trade rather than pretending it away.**

At group scope the same fixed point acquires an exact join to Paper A: `ln r = ln D − ln p_bar − F`, so **the group's surplus enters the doctrine-to-data ratio one nat for one nat**, and calibration becomes a geometry requirement rather than an attitude. That derivation is Chapter 47 §47.8 and it is the exact seam between Volume Two's two governing equations.

34.17 Sensitivity: the equation has unit elasticity in all three terms

Differentiating the log form:

```
ln M = ln D − ln P + ln E

dM/M = dD/D − dP/P + dE/E
```

Every elasticity is ±1 exactly. Verified numerically at D = P = E = 1: elasticity with respect to D is +1.000000, with respect to P is -0.999999 (finite-difference error), with respect to E is +1.000000.

A one per cent error anywhere produces a one per cent error in the moment. There is no term you can be sloppy about and no term whose error is amplified. That is unusual and it is a genuine virtue of the multiplicative form; a formula with a squared term or an exponent would have a sensitive parameter and a forgiving one, and would demand that measurement effort be allocated unevenly.

The asymmetry that does exist is between the numerator and the denominator at large deviations, because a ratio is not symmetric in its arguments:

perturbation	effect on M
D up 1%	x1.0100 (+1.00%)
D up 10%	x1.1000 (+10.00%)
D up 50%	x1.5000 (+50.00%)
P up 10%	x0.9091 (-9.09%)
P down 10%	x1.1111 (+11.11%)
P up 50%	x0.6667 (-33.33%)
P down 50%	x2.0000 (+100.00%)

Halving your intake doubles your moment. Halving your frame halves it. The two errors are not equally available or equally seductive, and the table shows why: the pathological direction is the one with the larger reward.

34.18 The moment is bounded: M <= ln N

Left as written, `M = (D/P).E` is unbounded, and §34.9 already flagged that a divergence to infinity is a modelling artifact. The bound comes from the register.

CANON_INTERNAL §1 requires that the moment, to serve as `S_t` in the storage law, be normalized into the surplus band `[0, ln N]`. The surplus law's ceiling is a theorem — Paper D, Corollary 2.2, `0 <= F(u) <= ln N` — and the normalization is a convention. Together:

```
0 <= M <= ln N              [THEOREM (Paper D Cor 2.2) + CONVENTION (normalization)]

  N = the number of categories the observer can distinguish
```

The richest possible moment is the logarithm of how many kinds of thing the perceiver can tell apart. This is not a poetic claim. It is Ashby's Law of Requisite Variety — only variety absorbs variety — appearing in first-person form, and it was not chosen; it follows from putting the moment into the surplus register at all.

Computed:

N (distinguishable categories)	ln N (max moment, nats)	in octaves (2F/ln 2)
2	0.693147	2.0000
3	1.098612	3.1699
5	1.609438	4.6439
7	1.945910	5.6147
10	2.302585	6.6439
12	2.484907	7.1699
20	2.995732	8.6439
50	3.912023	11.2877
100	4.605170	13.2877
1 000	6.907755	19.9316
10 000	9.210340	26.5754

And the brutal fact of Chapter 8, restated at this scope: **doubling your repertoire adds `ln 2 = 0.693147` nats to your maximum moment, no matter where you start.** `ln 100 − ln 50 = 0.693147`. `ln 4 − ln 2 = 0.693147`. Identical. Going from fifty distinguishable kinds of thing to a hundred buys precisely as much ceiling as going from two to four, and buys it at forty-eight times the cost.

The design consequence is severe and this manual states it plainly. **You cannot solve a distortion problem by learning more categories.** Raising N raises a ceiling you are almost certainly not touching. If your r is 2.1, you are manufacturing half of your moment, and adding vocabulary gives you more material to manufacture it from. **Fix r first. It is the multiplier; N is only the roof.**

34.19 A note on where the ceiling makes the divergence honest

With `M <= ln N` in force, Limit 1 reads correctly:

```
P -> 0   =>   M -> min(D.E/P, ln N)  ->  ln N
```

The dogmatic observer does not experience an infinite moment. They experience their maximum possible moment — the full extent of their categories, in high definition, sourced entirely from themselves. **Saturation, not divergence.** That is a much better description of the phenomenon and it is what the bounded model predicts, which is a small piece of evidence that the bound is real and not a patch.

34.20 The predictive-processing correspondence

CANON_INTERNAL §4 requires that this be stated, that it be stated carefully, and that it be stated as a correspondence between formalisms. All three.

Predictive-processing accounts in cognitive science (Friston; Clark; Hohwy) model perception as the weighting of top-down predictions against bottom-up sensory evidence, with the weights determined by the *precisions* — the inverse variances — assigned to each. In the simplest Gaussian case, with prior mean `mu` and precision `pi_D`, and sensory observation `y` with precision `pi_P`, the posterior mean is the precision-weighted average

```
mu_post = (pi_D . mu + pi_P . y) / (pi_D + pi_P)
```

Define the precision ratio `lambda = pi_D / pi_P` and the weight on the prior `w`:

```
w = lambda / (1 + lambda)          mu_post = w . mu + (1 − w) . y
```

Now the correspondence, and it is exact where it holds:

```
logit(w) = ln( w / (1 − w) ) = ln lambda
```

Verified:

lambda	w	1 - w	logit(w) = ln lambda
0.01	0.009901	0.990099	-4.605170
0.10	0.090909	0.909091	-2.302585
0.25	0.200000	0.800000	-1.386294
0.50	0.333333	0.666667	-0.693147
1.00	**0.500000**	**0.500000**	**0.000000**
2.00	0.666667	0.333333	+0.693147
4.00	0.800000	0.200000	+1.386294
10.00	0.909091	0.090909	+2.302585
100.00	0.990099	0.009901	+4.605170

The signed distortion of the Moment Equation, `ln r`, is the logit of the weight predictive processing assigns to the prior, when `r` is read as the precision ratio. The two formalisms place the same object at the same landmark: `lambda = 1` is `w = 1/2`, equal weight to prior and evidence, and it is `r = 1`, `M = E`, contact. Both put the invariant at the ratio one, and both measure departure from it in log units.

A worked instance, with a prior expectation of 100 and an observation of 60:

lambda	w	posterior	distance from the evidence
0.1	0.0909	63.636	+3.636
0.5	0.3333	73.333	+13.333
1.0	0.5000	80.000	+20.000
2.0	0.6667	86.667	+26.667
10.0	0.9091	96.364	+36.364

Now the caveats, and they are not decoration.

First, predictive processing is an active and contested research program, not settled fact. Its scope, its status as a theory of cortical function, and the strength of the free-energy formulation in particular are all under live dispute. Citing a correspondence with it is not citing a proof.

Second, structural similarity is not identity, and here the two objects differ in an important way. The Bayesian weight `w` is bounded in [0,1]: a precision-weighted average cannot land outside the interval between the prior and the evidence. The Moment Equation's `r` is unbounded above: `M` can exceed `E` without limit, up to the `ln N` ceiling. These are structurally different operations. **The Bayesian model interpolates; the Moment Equation scales.** The logit identity maps one onto the other cleanly at the level of the log-ratio parameter, and does not make the models the same model. An observer whose posterior sits between prior and evidence is not manufacturing content; an observer with `r = 4` is.

Third, the direction of inference runs only one way and it is the weak way. That two independently constructed formalisms put their invariant at the same place is interesting. It is not evidence that either is correct, and this manual does not use it as such. It is offered because CANON_INTERNAL §4 calls it genuinely striking and it is, and because a reader who knows the predictive-processing literature should be told immediately where the two frameworks touch and where they do not.

34.21 The join: the moment becomes character

The Moment Equation is not the end of the pipeline. It is the input stage. CANON_INTERNAL §8:

```
THE SELF–STORAGE LAW      Q_{t+1} = (1 − δ) Q_t + η M_t − C_t
```

which is the storage law of Paper E §3 with the moment substituted for the surplus input. **This is the join between the two formulas, and it is the architecture of Volume Two: perspective and perception produce the moment; the moment becomes character.** Chapter 40 develops it in full.

What belongs here is the consequence for calibration, and it is unpleasant.

Take a person with `η = 0.35`, `δ = 0.08`, `C = 0.20`, in an environment worth `E = 1.0` nat. The steady state of the self-storage law (Paper E, Theorem 3.5) is `Q_inf = (ηM − C)/δ`:

condition	r	M = rE	Q_inf
Under-framed	0.35	0.3529	**-0.9559**
Calibrated	1.05	1.0476	+2.0833
Over-framed	1.90	1.9000	**+5.8125**

Read the last row and then read it again. **The over-framed observer stores nearly three times as much as the calibrated one.** The storage law integrates the moment as lived, and it does not inspect it. There is no term in `Q_{t+1} = (1−δ)Q_t + ηM_t − C_t` that asks whether `M` was earned.

That is not a defect in the storage law. It is the storage law telling the truth about a real phenomenon: **dogma accumulates faster than knowledge, and it accumulates by exactly the same mechanism.** The confident, model-driven, low-intake observer builds a large, stable, internally consistent structure very efficiently. It is worth noting that the under-framed observer, by contrast, has a *negative* steady state — `ηM < C`, the regime Paper E Theorem 3.5 explicitly refuses to hide — which means the flooded novice is not slowly building; they are running down. The model says the flooded state is unsustainable and must be exited, and that is correct.

34.22 Veridical throughput: separating what was stored from what was earned

If the storage law cannot distinguish earned structure from manufactured structure, the engineer must, and the equation supplies the instrument.

> **VERIDICAL THROUGHPUT [DEFINITION].** The part of the moment that the environment actually supplied is
>
> ```
> V = min(M, E) = E . min(r, 1)
>
> confabulated fraction   = (M − V)/M = 1 − 1/r     for r > 1
> missed fraction         = (E − M)/E = 1 − r       for r < 1
> ```

Computed:

r	M/E	V/E	confabulated fraction	missed fraction
0.10	0.1000	0.1000	—	0.9000
0.25	0.2500	0.2500	—	0.7500
0.50	0.5000	0.5000	—	0.5000
0.75	0.7500	0.7500	—	0.2500
1.00	**1.0000**	**1.0000**	**—**	**—**
1.33	1.3300	1.0000	0.2481	—
2.00	2.0000	1.0000	0.5000	—
4.00	4.0000	1.0000	0.7500	—
10.00	10.0000	1.0000	0.9000	—

V is maximised exactly at the fixed point, and only at the fixed point. Over-framing does not increase what the environment gave you; it adds material that came from you and is stored as if it came from outside. Under-framing loses what was there. The fixed point is the unique maximum of veridical throughput, and that is the fourth and strongest sense in which D = P is the formula's centre: it is not merely where the moment is undistorted, it is where the earned content of the moment is largest.

Running both quantities through the storage law separates stored structure from stored error:

condition	Q_inf on M (as lived)	Q_inf on V (veridical)	stored error
Under-framed, r = 0.35	-0.9559	-0.9559	0.0000
Calibrated, r = 1.05	+2.0833	+1.8750	+0.2083
Over-framed, r = 1.90	+5.8125	+1.8750	**+3.9375**

The over-framed observer's structure is 68 per cent stored error at steady state. It will be defended as vigorously as the earned 32 per cent, because from inside there is no difference: **stored error and stored knowledge occupy the same variable.**

The one mechanism that eventually distinguishes them is cost. Structure that does not match the environment generates correction events, contradictions, failures and repairs — and those are `C`. **Over-framing shows up in the books not as bad `Q` but as rising `C`, on a lag.** That lag is why

the condition is stable for years and why the diagnostic instruments of Chapter 35 have to be run deliberately rather than waited for.

34.23 What the equation does not do

- It does not compute what is really there. E is exogenous and must be measured by an instrument outside the observer.
- It does not tell you whether a given moment is good. It tells you whether it is faithful.
- It does not separate D from P. FVP-5 is open on exactly this point: `(D, P, E)` and `(kD, kP, E)` make identical predictions. Only the ratio is identified, and any claim about the absolute magnitude of a person's frame or intake needs a second instrument.
- It does not apply to environments that respond to the observer. Where the observer's assertion changes E — a negotiation, a diagnosis that becomes a treatment, a forecast that moves a market — E is not exogenous and the formula as written is inapplicable. **That is a substantial class of real situations and the manual does not paper over it.** The reflexive case requires a coupled system, which is not developed in this volume.

34.24 FVP-6 in full: what would kill this formula

Following the Paper E §5.3 template.

Preregistration. Fix the operationalization of D, P and E (§34.25's count-based estimator or another declared one), the population, the item set, the number of occasions, and the analysis, in writing, before data collection.

Fit and freeze. Estimate `r` for each observer on a training block. Freeze every parameter.

Held-out prediction. Predict, on a held-out block, a quantity the formula constrains and that is not the quantity used to fit: the false-discovery rate, the off-model detection rate, and the ratio of asserted items to noticed items.

Four baselines, all of which must be beaten:

1. **Persistence** — predict the next block equals the last block.
2. **Direct-input** — predict from environment richness alone, ignoring the observer.
3. **Matched autoregressive** — an AR model on the observer's own past output with the same number of free parameters.
4. **Shuffled-input control** — the same model fitted to observer/environment pairings that have been randomly permuted. **This is the one most often omitted and it is the decisive one.** If the formula predicts as well on shuffled pairs as on true pairs, `r` is not carrying information about the observer and the formula has no content.

The killing result, stated in advance: if a frozen `r` fails to beat the shuffled-input control on held-out prediction, the Moment Equation is dead as a quantitative model and must be demoted to the heuristic ordering of §34.4. **A framework that cannot fail cannot be trusted, and this is exactly how this one fails.**

FVP-7, the reproducibility gate, requires the same operationalization run by different investigators on different populations, with the question: does `r` recover as a property of the observer, or does it turn out to be a property of the item set? If `r` for the same observer swings by more than the calibrated band width (0.10 nats) across equivalent item sets, `r` is not a parameter of

the person and the framework's psychological reading collapses even if its item-specific reading survives. **Chapter 54 argues that the Rasch model's separability property is the standard this gate should be held to, and this manual holds it to that standard.**

The Specification

S34.1 — Commensurability of D and P

D and P **shall** be measured in identical units, by procedures that differ only in which of the two they measure. **A ratio of two quantities that are not the same kind of quantity shall not be written as a fraction anywhere in a document governed by this manual.**

S34.2 — Ratio scale or demotion

Before any numerical use of `M = (D/P).E`, the analyst **shall** state the scale type of each term per Stevens' typology. Where any term is ordinal, **the formula shall be printed in its demoted heuristic form and no number shall be computed from it.**

S34.3 — E measured independently

E **shall** be measured by an instrument that does not depend on the observer whose moment is being modelled — a seeded item list, an independent panel, a ground-truth record. **An E inferred from the observer's own report makes the equation circular and its outputs void.**

S34.4 — Report r, not D and P

Because only the ratio is identified (FVP-5), reports **shall** give `r` and `ln r`. Separate values of D and P **shall** be reported only where a second, independent instrument fixes the scale, and **shall** carry a note that they are not identified by the moment measurement alone.

S34.5 — Report veridical throughput alongside the moment

Any report of M **shall** be accompanied by `V = E.min(r,1)` and by the confabulated or missed fraction. **Reporting a large moment without its veridical throughput is the single most misleading thing this instrument can be used to do.**

S34.6 — Distortion in log units

Distortion **shall** be reported as `ln r` (signed) and `|ln r|` (magnitude), never as a raw ratio or a percentage. See Chapter 35 for the band table and Chapter 36 for why.

S34.7 — Correction gain within the stability band

Any deliberate correction procedure **shall** target a gain `0 < a < 2`, and **shall** be inspected for `a >= 2` behavior — alternating overshoot of increasing amplitude — at least once per quarter for any observer or body under correction. **A correction regime that produces alternating extreme positions shall be stopped, not tuned.**

S34.8 — No calibration claim without a seeded audit

A claim that an observer, panel or institution is calibrated **shall** rest on a measurement against known items. **Self-assessed calibration shall be recorded in the survey as absent.**

S34.9 — Off-model items mandatory in every audit

Every seeded audit **shall** include items outside the frame under test, at a stated proportion of at least 20 per cent of seeded items. **An audit containing only on-model items cannot detect over-framing and shall not be used to certify calibration.** This is the highest-value requirement in the chapter.

S34.10 — Environment checked before the observer

Where M is low, E **shall** be measured before any intervention on the observer. **A calibrated observer reporting a poor environment is a correct instrument and shall not be treated as a perception problem.**

S34.11 — The ceiling declared

Any use of the normalized moment **shall** state the assumed N and therefore the ceiling `ln N`, and **shall** state that the normalization is a convention while the ceiling `0 <= F <= ln N` is a theorem (Paper D, Cor 2.2).

S34.12 — Reflexive environments excluded

Where the observer's assertion materially changes E, the analyst **shall** record the situation as out of scope for the Moment Equation as written, and **shall not** report an `r` for it.

S34.13 — Decline path

No calibration procedure in this manual **shall** be administered to a person who has not agreed to it, and a refusal **shall** be recorded as a refusal and not as a result. CANON §5.4 governs absolutely.

The Failure Modes

FM-34.1 — The Incommensurable Fraction

Description. D and P are measured in different units — years of experience over items noticed, confidence rating over minutes spent — and divided anyway.

Damages. Every downstream number. The output is a quantity with units nobody can name.

Diagnostic signature. Ask what `r = 1` would mean physically. If the answer is not a sentence about the same quantity appearing twice, FVP-1 has failed.

Fix. S34.1. Rebuild both measurements from one procedure. In the audit design of §34.25 both are *recorded items per condition present*, which is why that design works.

FM-34.2 — The Likert Fraction

Description. D and P come from five-point self-report scales and are divided.

Damages. FVP-2. The number produced has the form of a ratio and none of the content.

Diagnostic signature. The instrument has no true zero. Ask what a score of 0 would be and whether a 4 is genuinely twice a 2.

Fix. S34.2. Demote the formula in print, or build a ratio scale by magnitude estimation (Chapter 38) or by counting (Chapter 53).

FM-34.3 — Certainty Read as Calibration

Description. An observer reports high confidence and low confusion, and this is taken as evidence of good contact.

Damages. Detection of `P -> 0`. The condition is misdiagnosed as competence.

Diagnostic signature. **Falling surprise rate combined with a rising false-discovery rate.** The observer is being contradicted less and being wrong more. Also: time-per-case falling while off-model detection falls.

Fix. S34.8 and S34.9 — a seeded audit with off-model items. There is no soft version. Confidence is the symptom of the disease, so it cannot be the test.

FM-34.4 — The Empty Aperture Mistaken for Discipline

Description. An observer with a strong frame reads their own low intake as focus, and the organisation rewards it as decisiveness.

Damages. `P`, therefore `r`, therefore `V`.

Diagnostic signature. Off-model hit rate near zero with on-model hit rate near one. In the worked example of Chapter 35, the veteran inspector detects 94 per cent of expected conditions and 13 per cent of unexpected ones.

Fix. Rebuild intake specifically for off-model material: mandated fresh-eyes passes, rotation of inspectors across building types, an explicit "what is here that is not on the list" step with its own time allocation.

FM-34.5 — Frame Removal Prescribed as Clarity

Description. An under-performing observer is advised to "drop their assumptions", "come to it fresh", "stop overthinking". D is driven toward zero.

Damages. `D`, therefore `M`. Drives the observer into Limit 2.

Diagnostic signature. Intake high, output thin, time-per-case rising, the observer describing the work as overwhelming rather than as wrong.

Fix. This is an under-framing case and the treatment is the opposite of the advice given: supply structure. Checklists, worked exemplars, a taxonomy, an apprenticeship. **The instruction "have no expectations" is a prescription for `M = 0` and this manual does not issue it.**

FM-34.6 — The Environment Blamed on the Observer

Description. M is low. The observer is sent to training.

Damages. Wastes the intervention budget and leaves E unimproved; damages the observer's standing; often raises C.

Diagnostic signature. A seeded audit shows `|ln r| < 0.10`. The instrument is fine.

Fix. S34.10. Measure E. Chapter 12 and the whole of Volume One are the intervention.

FM-34.7 — Stored Error Defended as Experience

Description. A long-tenured observer or institution has accumulated large `Q` at high `r`. The structure is coherent, confident, and 60–70 per cent manufactured. Challenges to it are read as challenges to competence.

Damages. `Q` (as stored error), and `C` on a long lag as corrections accumulate.

Diagnostic signature. Compute `Q_inf` on M and on V separately (§34.22). A gap of more than about a third of `Q_inf` is the signature. Cross-check with the trend in `C`: **rising rework, rising dispute, rising correction cost, with no change in the environment, is the lagging indicator of a high-`r` history.**

Fix. Shedding stored order to regain headroom is legitimate and underused (Chapter 6's fourth lever). Combine with a correction regime at gain `a` in (0,1) — see Chapter 35 §35.16. Do not attempt deadbeat correction (`a = 1`) on a large `Q`; the cost is paid in one period and it is usually more than the system's admissibility condition allows.

FM-34.8 — The Divergent Corrector

Description. An observer or body responds to each contradiction by reversing position entirely, then reverses again on the next. Correction gain `a >= 2`.

Damages. `r`, with increasing amplitude. From `|ln r| = 0.81`, six iterations at `a = 2.3` reach `|ln r| = 3.92`, a factor of fifty from contact.

Diagnostic signature. Alternating sign of `ln r` with growing magnitude across successive audits. Rhetorically: "we were completely wrong about that" appearing at a regular cadence, each time about the previous correction.

Fix. S34.7. Reduce the gain. The stability condition `0 < a < 2` is the whole of the theory here, and the practical form of it is: **never move more than your evidence, and never move twice on the same evidence.**

FM-34.9 — The Unmeasured Environment

Description. E is estimated from the observer's own account, and `r` is computed from it.

Damages. Circularity. `r` comes out near 1 by construction, and the audit certifies calibration that was assumed.

Diagnostic signature. Ask where the ground truth came from. If the answer is the observer, the number is void.

Fix. S34.3. Seed the items. A defect list, a planted set of conditions, an independent adjudicated record.

FM-34.10 — The Reflexive Case Forced Into the Formula

Description. r is computed for a negotiator, a forecaster whose forecast moves the market, or a clinician whose diagnosis becomes the treatment. E is not exogenous.

Damages. Everything. The formula's central assumption is violated.

Diagnostic signature. Ask whether the environment would have been the same had the observer concluded otherwise. If not, out of scope.

Fix. S34.12. Record as out of scope. Do not report an r.

Worked Example

The seeded-defect review audit

A software group runs code review. One reviewer, R1, is respected, fast, and finds a great deal. The group wants to know whether R1's reviews are contact or manufacture. Nobody has ever asked this about a reviewer before, which is the point.

Step 1 — Build E. The team constructs a review packet by seeding a real, already-shipped module with 40 known defects. The seeding is done by two people who will not participate in the review, and adjudicated in advance so that E is fixed before any reviewer sees it. Of the 40:

- 30 are **on-model**: null checks, off-by-one, resource leaks, the categories R1's team lists in its review guide.
- 10 are **off-model**: a licence incompatibility, a timezone assumption, a comment that contradicts the code, an error message that leaks a path, and six others outside the guide.

Off-model items are 25 per cent of the total, satisfying S34.9. **Without them the audit cannot detect over-framing and would have certified R1 as excellent.**

Step 2 — Define D and P commensurably (S34.1). Both are *recorded items per condition present*:

```
   P_hat = (real conditions appearing anywhere in the reviewer's working notes) / (condition
   D_hat = (total items in the reviewer's submitted review)              / (conditions pres

        both  in  units  of  [recorded  items]  /  [conditions  present]   -
>  r = D_hat / P_hat is dimensionless
```

The working notes are the intake record: everything the reviewer wrote down while reading, including things they later dropped. The submitted review is the assertion record. **The two records are the same kind of object — items the reviewer wrote — which is what makes the fraction legal.**

Note the honest caveat that FVP-1 forces us to print: this treats a note and a review comment as the same unit of "recorded item". That is an assumption. It is defensible because both are

discrete written assertions about the module, and it is checkable — two independent counters should agree on the count to within a few per cent — but it is an assumption and it is the weakest joint in the measurement.

Step 3 — Run the audit, baseline.

```
conditions present                    40   (30 on-model, 10 off-model)
real conditions touched in the notes  29
items in the submitted review         50   (29 real, 21 not real)

P_hat = 29 / 40 = 0.725000
D_hat = 50 / 40 = 1.250000
r     = 1.250000 / 0.725000 = 1.724138
ln r  = +0.544727 nats
```

R1 is over-framed by a factor of 1.72. The moment R1 lives while reviewing is 1.72 times the module. By the Chapter 35 band table this is MARKED over-framing, one band below severe.

There is an identity worth noting here, because it makes the estimator far cheaper than it looks. Since `r = (calls/present)/(noticed/present) = calls/noticed`, and since in this audit every real condition that reached the notes also reached the review, we have

```
   r = 50/29 = 1.724138
confabulated fraction = 1 − 1/r = 1 − 29/50 = 0.4200 = false discovery rate exactly
```

When everything noticed is asserted, the confabulated fraction of the moment is exactly the false discovery rate. `r` is then computable from two counts and needs no ground-truth adjudication at all — though ground truth is still required to measure P against E and therefore to distinguish this case from an under-framed one.

Step 4 — Read the secondary indicators.

```
on-model hit rate     28/30 = 0.9333
off-model hit rate     1/10 = 0.1000
false discovery rate  21/50 = 0.4200
intake (P_hat)              0.7250
```

The signature of FM-34.4 is present in the open: 93 per cent detection of what the guide predicts, 10 per cent detection of what it does not. R1 is not reading the module. **R1 is checking the module against the guide, and the guide is also supplying 21 findings that are not there.**

Step 5 — Compute veridical throughput and stored error. With E normalized to 1.0 nat:

```
M = r . E = 1.7241 nats
V = E . min(r,1) = 1.0000 nats
confabulated fraction = 0.4200
```

With the team's measured personal parameters $\eta = 0.35$, $\delta = 0.08$, $C = 0.20$:

```
Q_inf on M (as lived)   = (0.35 . 1.7241 − 0.20)/0.08 = +5.043
Q_inf on V (veridical)  = (0.35 . 1.0000 − 0.20)/0.08 = +1.875
stored error            = +3.168
```

R1's accumulated review expertise, at steady state, is about 63 per cent stored error. This is not a statement about R1's character. **It is a statement about what a guide-driven review process**

does to anyone who runs it for long enough without an off-model channel, and R1 is the group's best reviewer precisely because R1 runs the process hardest.

Step 6 — Design the correction. The distortion is `ln r = +0.5447`. Two levers, and the specification requires both to be tried, because `r` can be moved from either end:

Lower D. Cap review comments. Require that each submitted comment cite the line in the working notes that generated it — an assertion with no intake behind it cannot be filed. This directly attacks the 21 unfounded items.

Raise P. Add a mandated fifteen-minute off-guide pass with its own note field: "what is in this module that the guide does not cover". Rotate a second reviewer from a different subsystem once a week (raising u across reviewers, Chapter 7, which is a separate benefit).

Target gain `a = 0.35`, well inside the stability band, giving the trajectory of §34.14: five review cycles to the calibrated band.

Step 7 — Re-audit after six cycles. Fresh packet, same design, 40 seeded conditions, 30 on-model and 10 off-model.

```
real conditions touched in the notes   31
items in the submitted review          32   (31 real, 1 not real)

P_hat = 31/40 = 0.775000
D_hat = 32/40 = 0.800000
r     = 1.032258
ln r  = +0.031749 nats        CALIBRATED
```

Metric	Baseline	After six cycles	Change
`r`	1.7241	1.0323	-40.1%
`|ln r|` (nats)	0.5447	0.0317	**17.2x reduction**
Intake `P_hat`	0.7250	0.7750	+6.9%
On-model hit rate	0.9333	0.9000	-3.6%
Off-model hit rate	0.1000	0.4000	**4.0x**
False discovery rate	0.4200	0.0312	-92.6%
True defects found	29/40 = 0.725	31/40 = 0.775	+6.9%
Confabulated fraction	0.4200	0.0312	-92.6%
`Q_inf` stored error	+3.168	+0.141	-95.5%

Read the on-model row honestly: it went down. R1 now finds 90 per cent of guide-listed defects instead of 93 per cent. If the group measures reviewer quality by guide-listed defect detection — and most groups do — this intervention will register as a small regression by the metric that is being watched, while total true defects found rose, false alarms fell by 93 per cent, and detection of the defects nobody was looking for quadrupled.

That is Goodhart's Law arriving on schedule, and it is the last thing this worked example has to teach. **The measure that made R1 look excellent was the measure that was making R1 over-framed, because it counted only the items the frame already contained.** The audit that found it works solely because 25 per cent of its items were outside the frame. S34.9 is in the

specification for this reason, and it is the requirement to fight for when someone proposes to simplify the audit.

Step 8 — What would falsify this. The group preregisters, for the next quarter: R1's frozen `r = 1.03` predicts a held-out false-discovery rate below 0.10 and an off-model hit rate above 0.30, and does so better than (i) persistence from the baseline audit, (ii) a model using packet difficulty alone, (iii) an autoregressive model on R1's past comment counts, and (iv) the same model fitted to reviewer/packet pairings randomly shuffled across the team. **If the shuffled control predicts as well, `r` is a property of the packet and not of R1, the instrument is measuring nothing about the reviewer, and this chapter's central estimator fails FVP-7 for this application.** The group has agreed in writing to publish that result internally if it occurs.

That commitment, and not the arithmetic, is what makes the audit an instrument rather than an opinion.

Chapter 35 takes the ratio `r` on its own, gives distortion its measure `|ln r|`, and turns over-framing and under-framing into nameable states with diagnostic protocols and correction procedures in both directions. Chapter 36 explains why the measure has to be a logarithm — and in doing so answers the question this volume was written around.

Chapter 35 . The Calibration Ratio

Chapter 34 established that the observer's entire contribution to the moment passes through one dimensionless number. `M = (D/P).E` has four symbols and three degrees of freedom, and the third one — the only one the observer controls — is the ratio.

This chapter takes that ratio on its own and does four things with it. It gives distortion a measure. It turns the two directions of distortion into named clinical states with signature panels you can run on a Tuesday. It supplies estimators that work with and without ground truth. And it gives correction procedures in both directions, with a stability condition, a half-life, and an honest account of when you should not correct at all.

The chapter's working claim is that over-framing and under-framing are not temperaments or character flaws but measurable operating points of an instrument, that both are correctable, and that neither can be detected from inside.

The Law

35.1 The definitions

THE CALIBRATION RATIO [DEFINITION].

```
r = D / P                    dimensionless, r > 0
```

The ratio of frame to intake. `r = 1` is contact: the fixed point of Chapter 34, at which `M = E`.

THE DISTORTION MEASURE [DEFINITION].

```
      signed distortion      ln r              in nats, range (-
inf, +inf), zero at contact
   distortion        Δ = |ln r|  in nats, range [0, +inf), zero at contact
```

Both are stated as definitions rather than postulates, because they define quantities rather than assert facts. **What is postulated is that these quantities are worth measuring and that they behave as this chapter says they do; that claim is falsifiable and §35.30 says how.**

35.2 Why the measure is a logarithm and not a ratio

Three independent reasons force the log, and they are worth having separately, because a reader who rejects one still has two.

Reason 1 — symmetry. On the raw ratio scale, over-framing has infinite room and under-framing is crushed into the interval (0,1). That is a coordinate artifact and it corrupts every summary statistic computed on it. Under the logarithm:

r	ln r	1/r	ln(1/r)
2	+0.693147	0.5000	-0.693147
3	+1.098612	0.3333	-1.098612
4	+1.386294	0.2500	-1.386294

r	ln r	1/r	ln(1/r)
10	+2.302585	0.1000	-2.302585

`ln(1/r) = −ln r` identically. **Seeing twice as much as is there and seeing half as much as is there are the same magnitude of error in opposite directions, and only the log coordinate says so.** On the raw scale, r = 2 sits 1.00 above contact and r = 0.5 sits 0.50 below it — the same fault priced at half.

Reason 2 — composition. Lenses compose multiplicatively and logs turn that into addition:

```
r_total = r_1 . r_2 . ... . r_n
ln r_total = ln r_1 + ln r_2 + ... + ln r_n
```

Distortion budgets are additive in log units and only in log units. §35.19 builds a pipeline budget on this and it is not constructible any other way.

Reason 3 — register. CANON_INTERNAL §1: surplus, pitch, sensation and the moment are all logarithmic, and every conversion between them is linear in log space with exact constants. Reporting distortion in nats puts it in the same register as `F = ln N_eff`, as `p = log2 f`, and as Fechner's `S = k ln(I/I_0)`. Chapter 36 argues that this is not a convenience of notation but the register perception itself operates in, and that argument is the deeper form of Reason 1.

35.3 The band table

Distortion is a continuum and clinical work needs bands. The bands below are **a convention of this manual**, chosen for round numbers in nats and for correspondence with recognisable states. They are not measured constants and must not be reported as if they were.

`BAND TABLE — DISTORTION  [CONVENTION of this manual]`

Band	Δ = |ln r| (nats)	over-framed r	under-framed r	Reading
CALIBRATED	0.00 – 0.10	1.000 – 1.105	0.905 – 1.000	Contact. Within ±10.5 per cent of the environment.
MILD	0.10 – 0.35	1.105 – 1.419	0.705 – 0.905	Detectable, tolerable, worth trending.
MARKED	0.35 – 0.69	1.419 – 1.994	0.502 – 0.705	The moment is meaningfully not the environment. Correct.
SEVERE	0.69 – 1.61	1.994 – 5.003	0.200 – 0.502	Factor of two to five. Decisions from this reading are unsafe.

Band	Δ = |ln r| (nats)	over-framed r	under-framed r	Reading
EXTREME	1.61 and above	5.003 and above	0.200 and below	Factor of five or more. The observer is not in contact.

The band edges are `ln 2 = 0.693147` at the marked/severe boundary and `ln 5 = 1.609438` at the severe/extreme boundary — that is, **the severe band is exactly "wrong by a factor between two and five", in either direction, and the two directions occupy the same band width.** That is the symmetry of Reason 1 doing visible work.

35.4 Aggregation: the arithmetic mean of r is a lie

An observer measured on several occasions produces several values of r. The summary statistic matters enormously and the obvious one is wrong.

occasions	arithmetic mean r	geometric mean r	mean ln r	**mean |ln r|**
{0.5, 2.0}	1.250000	1.000000	0.000000	**0.693147**
{0.4, 0.8, 1.5, 2.5}	1.300000	1.046635	+0.045580	**0.615298**
{1.2, 1.25, 1.15, 1.3}	1.225000	1.223723	+0.201898	**0.201898**

Read row one. An observer who over-framed by a factor of two on one occasion and under-framed by a factor of two on another has an arithmetic mean r of 1.25, which reports mild over-framing that never occurred. Their geometric mean is exactly 1.000, which reports perfect calibration that also never occurred. **Their mean distortion is 0.693 nats — severe — and only the third statistic tells the truth.**

By Jensen's inequality the arithmetic mean of r is always at least the geometric mean, with equality only when every r is identical, so **the arithmetic mean of a calibration ratio always reports over-framing that may not be there.** Row three shows the degenerate case where all three agree: a steady, one-directional bias.

The rule, and it is short:

> **AGGREGATION RULE.** Report **mean `|ln r|`** as the distortion of an observer across occasions. Report **mean `ln r`** as the bias — the direction they lean. Report the geometric mean of r only as a back-transform of the second. **Never report the arithmetic mean of r.**

The pair (mean `ln r`, mean `|ln r|`) is the whole summary: a bias and a spread. An observer with bias +0.20 and distortion 0.20 leans one way consistently. An observer with bias 0.00 and distortion 0.69 is unstable in both directions and is a harder case, because they will pass any audit that averages.

35.5 FVP table for the calibration ratio

`FVP TABLE — THE CALIBRATION RATIO  r = D / P`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units; both sides; ratios dimensionless.	`[D] = [P]` by requirement S34.1; `r` is dimensionless. Inherits the commensurability condition and fails with it if that condition is violated.	**PASS**, conditional on S34.1
FVP-2 SCALE TYPE	Division requires ratio scale with true zero.	Requires D and P ratio-scaled on a **common** scale. The count-based estimator E1 (§35.13) supplies a zero and a unit *within one item set*, but not the cross-set linkage that would make the magnitude of `r` a measurement rather than a tally; rating-scale estimators supply neither. **Chapters 52 and 53 therefore return DEMOTE for the magnitude of `r` as the instrument stands.** Route out: **Ch 54 Thm 54.3** — Rasch logit difference, given badged Postulate SO1, at ~123 items per term for a half-logit distortion.	**DEMOTE** — direction and ordering stand, magnitude does not

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term -> 0, -> infinity.	`r –> 0`: total under-framing, moment vanishes (Ch 34 Limit 2). `r –> infinity`: total over-framing, moment saturates at `ln N` (Ch 34 §34.19). Both name real states, §35.9 and §35.10. `r` cannot be negative: D and P are counts or precisions, both non-negative.	**PASS**
FVP-4 FIXED POINTS	Solve for the invariant.	`r* = 1`. Unique. Under the correction map it is attracting for gain `0 < a < 2`, neutrally stable at `a = 0`, divergent at `a >= 2`.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable from obtainable data?	**Yes, and this is the strong gate for `r`.** Under estimator E1, `r = (items asserted)/(items noticed)` — two counts, no adjudication required for the ratio itself. Ground truth is needed only to distinguish the direction of a low `r` from a poor environment.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	Preregister, freeze, held-out, four baselines including shuffled-input.	Protocol at §35.30. Killing result stated: if frozen `r` fails to beat the shuffled-input control at predicting held-out false-discovery and off-model detection rates, `r` carries no information about the observer.	**NOT RUN — protocol stated**
FVP-7 REPRODUCIBILITY	Cross-individual, cross-site recovery.	Unknown. The decisive question is whether `r` is a property of the observer or of the item set. **If `r` for one observer varies across equivalent item sets by more than the calibrated band width of 0.10 nats, it is not a person parameter.** Chapter 54's separability standard applies.	**NOT RUN — the deciding gate**

35.6 FVP table for the distortion measure

`FVP TABLE — DISTORTION  Δ = |ln r|`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units.	`r` dimensionless, `ln r` in nats, Δ in nats. Convertible to bits (x1.442695) or octaves (x2.885390) by the Chapter 32 keys.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	What operations are legal on Δ?	Δ *would be* on a ratio scale with a true zero at contact — 0.40 nats genuinely twice 0.20 nats, two applications of the same multiplicative error — **and this is corrected here: that property is inherited from `r`, and `r`'s magnitude is demoted above.** Until D and P are co-calibrated, Δ is an **ordinal** quantity: it orders observers and occasions by how far they sit from contact, and it does not license the statement that one distortion is twice another. Chain addition (§35.19) is therefore a structural argument, not an arithmetic one, until Thm 54.3's instrument exists.	**DEMOTE to ordinal** — becomes ratio-scaled exactly when Δ is read as a Rasch logit difference (Ch 54)
FVP-3 LIMITING CASES	Δ -> 0 and Δ -> infinity.	`Δ = 0` iff `r = 1` iff `M = E`. `Δ –> infinity` requires `r –> 0` or `r –> infinity`, both bounded in practice by the `ln N` ceiling and by the finiteness of counts.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	Invariant.	`Δ = 0`, the same fixed point, expressed as a minimum rather than an equality. Δ is a **Lyapunov function for the correction map**: it decreases monotonically under any gain `0 < a < 1` and this is what proves the correction converges.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	Yes, from `r`. Note that Δ discards the sign and therefore **is not sufficient on its own to choose a correction** — §35.16 and §35.17 apply opposite procedures. Always report signed `ln r` alongside.	**PASS**, with the sign-loss warning
FVP-6 FALSIFICATION	As for `r`.	Inherits. Additional specific test: **does Δ predict downstream correction cost?** The framework claims high Δ raises `C` on a lag; a Δ that does not predict rework, dispute or repair costs is a Δ that is measuring nothing that matters.	**NOT RUN — protocol stated, and this is the highest-value single test in the volume**
FVP-7 REPRODUCIBILITY	As for `r`.	Inherits.	**NOT RUN**

35.6.1 What the demotion means for everything that follows

Both tables above now read DEMOTE on gate 2, and that verdict governs every number printed in the rest of this chapter. It is stated here, once, in full, so that no reader has to reconstruct it from a table cell.

> **`Delta` IS AN ORDINAL QUANTITY PENDING INSTRUMENTATION.** Its *sign* is a finding: an observer at `ln r > 0` is over-framed and an observer at `ln r < 0` is under-framed, and the two require opposite corrections (§35.15, §35.16). Its *rank order* is

a finding: an observer whose distortion falls between two audits has moved toward contact. **Its magnitude is not a finding.** "0.2851 nats", "1.33", "987 cents" — every such figure in this chapter and in Chapter 45 F-02 is an illustrative construction demonstrating the arithmetic of the measure, not a measurement of a person. The direction stands. The magnitude does not.

The reason is the one Chapter 53 works out in detail: D and P are not yet instrumented on a common ratio scale with a shared true zero, and division without a true zero produces a number with no referent. Counting recorded items gets you a zero and a unit inside one item set; it does not link two item sets, and without linkage the ratio moves when the items move.

The route out is constructive and it is already specified. Chapter 54, Theorem 54.3, given badged Postulate SO1: measure D and P as Rasch person-measures on one common linked scale, and `Delta = |ln(D/P)| = |theta_D − theta_P|` becomes ratio-scaled with a true zero at calibration. Rasch separation is what supplies the zero. The price is in §54.12: **about 123 items per term to establish a half-logit distortion at 95 per cent**, and about 31 items for a full logit. A twelve-item pair of scales resolves nothing below roughly 1.6 logits.

This is the protocol working. The chapter's own headline measure has been demoted by the chapter's own gate, the surviving content has been named exactly, and the study that would restore the magnitude has been priced in items. That sequence — demote, delimit, cost, specify — is what the Formula Validation Protocol is for, and it is worth more than a confident number would be.

The Engineering

35.7 The two states, stated once

```
r > 1    OVER-FRAMED     M > E    projection, prejudice, seeing the expected
 r = 1      CALIBRATED       M = E     contact
r < 1    UNDER-FRAMED    M < E    flooding, overwhelm, cannot make it mean anything
```

Neither state feels like an error from inside, and the two feel like opposite things. Over-framing feels like clarity, competence, speed, and confidence. Under-framing feels like exhaustion, fog, and inadequacy. **The pleasant one is the dangerous one**, and that asymmetry organises everything below: over-framing has to be hunted, under-framing announces itself.

35.8 OVER-FRAMED (r > 1) — the clinical description

The observer's assertions exceed their intake. Structure is being supplied from the frame and recorded as though it came from the environment.

Cardinal features.

Feature	What you see
High on-model detection	Finds what the model predicts, reliably and fast.
Low off-model detection	Misses what the model does not contain, almost completely.

Feature	What you see
High false-discovery rate	Asserts things that are not there, in the model's shape.
Falling time per case	Speed increases as intake falls, and is rewarded.
Falling surprise rate	Being contradicted becomes rare. Read as mastery.
High confidence	Subjectively: clarity.
Low revision rate	Positions do not change on new information.

Subtypes, which differ in what to do about them.

The doctrinal. The frame is institutional and enforced. Deviating from it is a career cost. The individual's `r` is a property of the organisation, and individual correction will fail. Treat at the level of Chapter 48.

The expertise decay. The frame was earned, was once accurate, and has outlived the environment that validated it. This is the commonest and most sympathetic form: **competence, correctly acquired, becoming distortion by the passage of time in the environment rather than by any error of the observer.** The worked example is a case of it.

The prejudicial. The frame is a category applied to persons in advance of evidence about them. Structurally identical to the others; the ethical weight is different and this manual does not pretend the arithmetic settles the ethics. What the arithmetic does say is that the mechanism is `P` falling, not `D` being wicked, and that the correction is intake.

The apophenic. The intake carries no signal — noise, sparse data, a genuinely ambiguous situation — and the frame supplies all of the structure. Distinguished from the others by the environment: `E` is thin, so any `D` at all produces a large `r`. **The correction here is often to withhold assertion rather than to change the frame**, and §35.24 covers that.

35.9 UNDER-FRAMED (r < 1) — the clinical description

The observer's intake exceeds their assertions. Material is arriving and not being made into anything.

Cardinal features.

Feature	What you see
High or adequate intake	Notices a great deal. The notes are long.
Low assertion rate	The output is thin relative to the notes.
Rising time per case	Everything takes longer, and the extra time buys little.
High surprise rate	Frequently caught out; the world keeps being unexpected.
Low confidence, high revision	Positions change on weak evidence.
Missed material	Not because it was not seen — because it was not concluded from.

Feature	What you see
Subjective report	Overwhelm; “I can’t make it mean anything.”

Subtypes.

The novice. The frame has not been built yet. Time and structure fix it. This is the benign case and the one everyone assumes they are looking at.

The overloaded. The frame exists and is adequate but the intake rate exceeds the rate at which it can be applied. **Distinguished from the novice by the observer’s performance on a smaller sample: give them a third of the material and see whether the assertions come.** If they do, this is a load problem, and the correction is on `P` or on the schedule, not on `D`.

The deframed. The observer has been instructed to drop their assumptions, come to it fresh, stop overthinking. `D` has been deliberately reduced by advice. This is FM-34.5 and it is an iatrogenic case — **the intervention caused the condition.**

The fragmented. The frame exists but is not retrievable under the conditions of use: distributed across notes nobody consults, or held by a person not in the room, or dependent on a system that is down. `D` is high in principle and zero in practice. The correction is availability, not training.

35.10 The differential diagnosis

The single instrument that separates the states, and the one to run first:

```
DIFFERENTIAL — the four conditions that all present as "poor performance"
```

Observation	Over-framed	Under-framed	Poor environment	Overloaded
`r` from seeded audit	> 1.105	< 0.905	~ 1	< 0.905
Intake `P`	Low	High	Normal	High
Assertion `D`	High	Low	Normal	Low
False discovery rate	High	Low	Low	Low
Off-model hit rate	Very low	Moderate	Normal	Low
Time per case	Falling	Rising	Normal	Rising
Surprise rate	Low	High	Normal	High
Performance on a **reduced sample**	Unchanged	Unchanged	Unchanged	**Improves markedly**
Independent measurement of `E`	Rich	Rich	**Thin**	Rich
Primary correction	Raise `P`, cap `D`	Raise `D`	Rebuild `E`	Cut load

Two rows in that table are the ones nobody runs and both are cheap. **The reduced-sample test separates the overloaded from the under-framed in a single session**, and they need opposite treatments: the overloaded observer given training gets worse, because training adds to the load. **The independent measurement of E separates every internal diagnosis from the case where the instrument is fine and the room is empty** (S34.10), and it is the first thing to check because it is the only diagnosis whose correction is outside the person.

35.11 Why you cannot detect your own r

The states are not introspectively distinguishable from their causes, and the reason is structural rather than psychological.

At `r > 1`, the internal report is: the situation is clear, I know what this is, this is going quickly. Every one of those reports is *also* what contact feels like. There is no experience attached to the twenty-one findings that were not there, because the observer does not experience them as manufactured — **`M` is the moment as lived, and a manufactured moment is lived exactly as a found one is.** The only internal signal that would discriminate is surprise, and surprise is precisely the quantity that falls as `P` falls.

At `r < 1`, the internal report is: I am not good enough at this. That report is indistinguishable from the report a calibrated observer would make in a genuinely hard environment, and it is also indistinguishable from the report of an overloaded observer.

Therefore: calibration cannot be self-assessed, and this manual does not accept self-assessment as evidence for it (S34.8). The instruments below are all external. That is not a comment on honesty; it is a consequence of the fact that the quantity being estimated is the very thing doing the estimating.

35.12 Estimating r with ground truth: the seeded audit

The reference method, specified in Chapter 34 §34.25 and used again in this chapter's worked example.

```
E    : a set of conditions planted or adjudicated in advance, independent of the observer
P_hat : (real conditions appearing in the observer's intake record) / (conditions preser
D_hat : (total items in the observer's assertion record)        / (conditions present)
  r      = D_hat / P_hat
```

Requirements: at least 20 per cent off-model items (S34.9), an intake record distinct from the assertion record, and ground truth from a source that is not the observer (S34.3).

Cost: one packet, one session, two counts. This is not an expensive instrument. It is a rarely used one.

35.13 Estimating r without ground truth: five behavioral estimators

Ground truth is not always available. Five estimators that are not, each with what it buys and what it is wrong about. **All five are proposed by this manual, none is validated, and each is badged.**

E1 — THE ASSERTION-TO-NOTICE RATIO [ESTIMATOR, unvalidated].

```
r_hat = (items asserted) / (items noticed)
```

From the algebra of §34.25: since both `D_hat` and `P_hat` divide by the same `E`, the ratio is just the two raw counts, **and the environment cancels.** This is the cheapest estimator in the manual and requires only that intake and assertion be separately recorded. *Biased by:* whether a note and an assertion are really the same unit. Requires S34.1 to be argued explicitly for the particular records used. *Blind to:* the case where both counts are low because the environment is thin.

E2 — THE SURPRISE RATE [ESTIMATOR, unvalidated].

```
surprise rate = (occasions on which the observer's expectation was contradicted) / (
```

Falls monotonically with `r` in the over-framed direction and rises with it in the under-framed direction. **Its value is as a trend, not a level:** a falling surprise rate in a stable environment is the earliest available indicator of drift toward over-framing, and it is available from ordinary records. *Biased by:* changes in the environment, which move it for reasons that have nothing to do with the observer. Requires an independently stable `E`.

E3 — THE REVISION RATE [ESTIMATOR, unvalidated].

```
revision rate = (positions changed after new evidence) / (occasions of new evidence)
```

Near zero at high `r`; high and unstable at low `r`. **Combined with E2 it separates the directions**: low surprise with low revision is over-framing; high surprise with high revision is under-framing; high surprise with low revision is an observer who is being contradicted and not updating, which is over-framing under pressure and is the worst of the four combinations.

E4 — TIME TO COMMIT [ESTIMATOR, unvalidated]. Time from first contact to first firm assertion. Short and falling at high `r`; long and rising at low `r`. *Biased by:* incentives on speed, which are usually present and usually push `r` up. **In any organisation that measures throughput, expect a systematic push toward over-framing**, and read E4 accordingly.

E5 — THE OFF-MODEL RATIO [ESTIMATOR, unvalidated].

```
off-model ratio = (assertions outside the standing model) / (assertions total)
```

Requires only a written model to compare against. A ratio near zero over many occasions is evidence of over-framing **or** of an environment that genuinely contains nothing off-model, and those must be distinguished by a seeded item or by rotation.

Use at least two, from different rows. E1 and E5 together are the cheapest useful pair: one gives magnitude, the other gives the over-framing direction, and both come out of records that already exist.

35.14 The correction map

Correction is a first-order recursion in log space, introduced in Chapter 34 §34.14 and used here as procedure:

```
ln r_{t+1} = (1 − a) . ln r_t            [POSTULATE — the correction map]
```

```
a  = correction gain: the fraction of the observed discrepancy acted on per cycle
 fixed point ln r* = 0
 stability: converges iff |1 − a| < 1, i.e. 0 < a < 2
```

Distortion half-lives, computed:

gain a	pole |1-a|	half-life (cycles)	cycles to one tenth
0.05	0.95	13.513	44.891
0.10	0.90	6.579	21.854
0.20	0.80	3.106	10.319
0.35	0.65	1.609	5.345
0.50	0.50	1.000	3.322
0.65	0.35	0.660	2.193
0.80	0.20	0.431	1.431

Δ is halved every `ln 2 / (−ln|1−a|)` cycles, which is the same half-life formula Chapter 3 gives for the decay of stored order under δ, with `|1−a|` in the role of `(1−δ)`. It is the same circuit. Calibration is a leaky integrator whose leak is the correction procedure, and the correction procedure is the only leak it has: **with no correction, distortion does not decay. There is no natural forgetting of a bad frame.**

That last sentence is the practical heart of the chapter. Chapter 3 spends its length on how to engineer δ down because stored order decays without maintenance. Distortion is the opposite: it persists without maintenance, because the storage law integrates it as readily as anything else (Chapter 34 §34.21). **You do not drift back into calibration. You are corrected into it or you stay where you are.**

35.15 Correcting over-framing (r > 1)

Two levers. The specification requires both to be considered, because the same Δ can be reached from either end and the cheaper lever differs by case.

Lever A — lower D. Make assertion cost something.

Procedure	Mechanism	Typical effect on r
Citation requirement	Every assertion must name the intake record entry that generated it. An assertion with no intake behind it cannot be filed.	Direct cut to the numerator; in the Chapter 34 worked example this removed 20 of 21 unfounded items.
Assertion cap	A hard limit on items per case, forcing prioritisation.	Cuts D; risks cutting true findings too. Use with a measured audit, not alone.
Explicit alternative	Require one written alternative reading of the same evidence before committing.	Reduces the frame's monopoly on the intake.

Procedure	Mechanism	Typical effect on r
Falsifier statement	Require, per assertion, the observation that would show it wrong.	Converts unfalsifiable frame-output into a testable claim; several will not survive the writing.

Lever B — raise P. Buy intake back.

Procedure	Mechanism	Typical effect on r
The off-model pass	A separately timed pass whose only question is "what is here that the model does not cover", with its own note field.	Raises the denominator and directly attacks the low off-model hit rate. Highest-value single procedure.
Rotation	Move the observer across contexts so the standing model stops matching.	Raises P by making the frame visibly insufficient. Also raises u across the body (Chapter 9).
Fresh eyes	A second observer with a different lineage, reporting independently before comparison.	Reveals the off-model material the first observer's frame filtered. Requires Chapter 22's order-of-operations rule: **independent record before discussion, or the second observer inherits the first one's frame.**
Slowing the case	Restore time removed by throughput incentives.	Raises P; unpopular; often the only thing that works when E4 shows an incentive-driven r.
Seeded items in live work	Plant known off-model conditions in real caseloads at low rate.	Maintains intake by making it consequential. Requires consent and governance.

Lever B is usually the better one and is almost never the one chosen. Lowering D is visible, fast and satisfying; raising P costs time, which is the resource the organisation is already short of. But cutting D alone moves `r` toward 1 while leaving the observer's contact with the environment exactly where it was, and the veridical throughput `V = E.min(r,1)` does not improve at all until P improves. **Cutting D fixes the ratio; raising P fixes the reading.**

35.16 Correcting under-framing (r < 1)

Lever C — raise D. Supply structure.

Procedure	Mechanism
Worked exemplars	Show completed cases with the reasoning visible. Builds the pattern library faster than instruction.
A written taxonomy	Give the categories names. This is literally raising N and it raises the ceiling as well as the frame.
Checklist with a decision rule	Not merely a list of things to look at, but a rule for what each observation implies. **A checklist without a decision rule raises P and not D and will make under-framing worse.**
Apprenticeship	High-η structure transfer (Chapter 4). Expensive, slow, and the most reliable.
Forced commitment	Require a written conclusion before leaving the case, even a provisional one. Converts intake into assertion by fiat and generates the error signal that correction needs.

Lever D — lower P. This is legitimate and is usually dismissed.

For the overloaded observer, whose frame is adequate and whose intake exceeds their capacity to apply it, the correction is to reduce the intake: smaller batches, fewer simultaneous cases, a filter upstream, a longer interval. **Turning down the input is a correction, not a retreat, and the reduced-sample test of §35.10 is what licenses it.**

The engineering statement: r is a ratio, and a ratio can be corrected from either term. An organisation that only ever responds to under-performance with training is only ever pulling one lever, and it is pulling the expensive one.

35.17 Choosing the gain

The gain a is a design parameter and the table of §35.14 is its specification.

- a **below about 0.10** is under-correction. It converges — twenty-two cycles to reduce distortion tenfold — but slowly enough that the environment may change first.
- a **between 0.20 and 0.50** is the working band. Half-lives of three cycles down to one.
- a **near 1.00** is deadbeat correction: one step to zero, in theory. In practice it means discarding the entire existing frame in one move, and the cost of doing so appears immediately in the storage law's admissibility condition C_t <= (1−δ)Q_t + ηM_t. **Deadbeat correction of a large Q violates admissibility and produces the collapse it was meant to prevent.**
- a **between 1 and 2** oscillates and converges. Wasteful but safe.
- a **at or above 2** diverges. FM-34.8. From Δ = 0.81, six cycles at a = 2.3 reach Δ = 3.92 — a factor of fifty from contact, arrived at by correcting hard every single time.

The design rule is one line: never move more than your evidence, and never move twice on the same evidence. Both halves are needed. The first bounds a below 2 on a single step; the second stops the same discrepancy being counted repeatedly, which is the mechanism by which a well-intentioned corrector reaches an effective gain above 2 without ever deciding to.

35.18 Composition, and the compensating-distortion trap

Distortion composes along a chain of observers:

```
r_total = r_1 . r_2 . ... . r_n          ln r_total = Σ ln r_i
```

The consequence that matters:

chain	net r	net Δ	total Δ injected
2.5 then 0.4	1.0000	0.000000	**1.832581**
0.5 then 2.0	1.0000	0.000000	**1.386294**
1.6 then 0.625	1.0000	0.000000	**0.940007**

A severely over-framed observer feeding a severely under-framed one produces a perfectly calibrated output and a completely broken system. The net distortion is zero. The total distortion present is 1.83 nats. The end-to-end measurement certifies contact.

This is not a curiosity. It is the standard structure of a hierarchy in which an enthusiastic analyst reports to a sceptical manager, and it is *stable*, because the end-to-end numbers look good and each party's correction of the other maintains the balance. It fails the moment either party is replaced, and it fails without warning, because the failure is a change in a quantity nobody was measuring.

> **DECOMPOSITION RULE.** Distortion **shall** be measured per stage, not end to end. Report both `|ln r_total|` (net) and `Σ|ln r_i|` (injected). **A system in which the injected distortion materially exceeds the net distortion is a compensating system and shall be recorded as fragile.**

35.19 The distortion budget

Budgeting in log units, along a four-stage reporting pipeline:

Stage	r	ln r	cumulative r	cumulative ln r	injected Δ so far
Field intake	0.85	-0.162519	0.850000	-0.162519	0.162519
Write-up	1.15	+0.139762	0.977500	-0.022757	0.302281
Supervisor review	1.35	+0.300105	1.319625	+0.277348	0.602385
Client summary	1.60	+0.470004	2.111400	+0.747351	1.072389

End-to-end `r = 2.111`, net distortion 0.747 nats (marked, edging severe), total injected 1.072 nats.

Two readings. First, **no single stage is worse than mild-to-marked, and the end product is severely over-framed**, because distortion accumulates additively in log units and each stage's contribution is small enough to survive its own review. Second, the pipeline briefly passes through near-calibration at stage two (cumulative `ln r = −0.0228`), which means a measurement taken

at the write-up would certify the pipeline as calibrated. **Where you measure decides what you find, and a single measurement point in a multi-stage pipeline is not a measurement of the pipeline.**

A budget is set the same way an error budget is set in any engineering discipline: allocate a total, divide it across stages, measure per stage, and require that the sum stay under the total. A defensible starting allocation for a four-stage pipeline with a target of "the client receives a mildly distorted reading" is a total budget of 0.35 nats, which is 0.0875 nats per stage — that is, **every stage must stay inside ±9 per cent.** That is demanding, and stating it as a number is the point: most pipelines have never been told what their per-stage tolerance is, and consequently do not have one.

35.20 When r should not be 1

The manual is required to say this and the chapter would be dishonest without it.

Calibration is not always the objective, because the moment is not always the product. The Moment Equation says `r = 1` maximises fidelity and veridical throughput. It does not say `r = 1` minimises loss, and where the costs of the two error types differ, the loss-minimising operating point is not the fidelity-maximising one.

Take the worked example's three operating points and price them under three cost structures:

Operating point	r	misses	false alarms	cost at 1:1	cost at 10:1 (miss:alarm)	cost at 1:3 (miss:alarm)
Novice	0.353	41	3	44	413	50
Trained	1.048	10	4	**14**	**104**	**22**
Veteran	1.900	20	27	47	227	101

In this instance the calibrated point wins under all three cost structures, which is the comfortable case. It does not always. **Where a miss is catastrophic and a false alarm is cheap, a deliberate `r > 1` is a rational policy**: screening, triage, safety interlocks, and early warning systems are all built to over-assert on purpose, and their false-discovery rates are a design choice rather than a fault. Conversely, **where assertion is expensive and irreversible — a public accusation, a permanent record, a decision that cannot be unwound — a deliberate `r < 1` is rational**: withhold conclusion, keep taking intake, decline to make it mean something yet.

The distinction the manual insists on:

> **Calibration is a property of the instrument. The operating point is a policy.** An observer may be perfectly calibrated and still choose to assert above or below their intake because the loss function says so. **What is never legitimate is an operating point that was not chosen — a policy-level `r` that is actually an uncorrected instrument error, defended after the fact as prudence or as decisiveness.**

The test is whether the observer can state the loss function. A screener who says "we run hot because a miss costs a hundred times a false alarm" has a policy. A reviewer who says "I have a good eye for these" has an uncorrected `r`, and the two are indistinguishable from the outside until you ask.

Two structural reasons, added here as corrections to any reading of this chapter that treats `r = 1` as a simple goal.

First: calibration converges to zero headroom, exactly. At `r = 1` the input to the storage law sits at its ceiling, `S = ln N`. Paper E Theorem 3.5 then gives `Q_inf = (eta ln N − C)/delta`, and the master gauge reads

```
h = eta ln N − delta Q_inf − C = eta ln N − (eta ln N − C) − C = 0
```

by cancellation, with `delta` and `C` dropping out. Perfect contact drives the input to the maximum, the maximum input drives the store to its maximum, and the maximum store spends the entire budget on maintenance. Chapter 34 §34.16 derives this at personal scope and Chapter 47 §47.6 at institutional scope. **`r = 1` maximises fidelity and veridical throughput; it does not leave room to build.** An observer who wants headroom at a settled state must shed `Q` deliberately (Chapter 40 §40.9) or hold a declared `r != 1`.

Second, and sharper: an observer optimising raw stored order is driven away from `r = 1` altogether. Chapter 34 §34.21 establishes the canon that over-framing *increases* stored order — the storage law integrates the moment as lived and never asks whether it was earned — and Chapter 40 §40.2 turns that into the two-track ledger: `Q` against veridical `Q^V`, with stored error `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` and correction load `C_t = c_0 + kappa Z_t`. Because `V = E.min(r,1)` is flat above the fixed point while `M = rE` keeps climbing, the veridical track stops improving at `r = 1` and the raw track does not. Computed at `eta = 0.35`, `delta = 0.08`, `c_0 = 0.20`, `E = 1.0`:

r	M	V	Z_inf	Q_inf (raw track)	Q_inf^V (veridical track)
0.50	0.5000	0.5000	0.0000	-0.3125	-0.3125
1.00	1.0000	1.0000	0.0000	**1.8750**	**1.8750**
1.50	1.5000	1.0000	2.1875	3.5156	1.8750
3.00	3.0000	1.0000	8.7500	8.4375	1.8750
5.00	5.0000	1.0000	17.5000	15.0000	1.8750

(Raw column computed with correction load at `kappa = 0.02`; the veridical column is flat above the fixed point by construction. Illustrative construction, per §35.6.1 — the structure is the finding, not the magnitudes.) **Dogma banks faster than knowledge, and the bill arrives on a lag as correction load.** Differentiating for `r > 1` gives the exact condition under which the bill ever catches up:

```
dQ_inf/dM = eta (1 − kappa/delta) / delta      −>   negative iff  kappa > delta
```

The knife edge is exactly `kappa = delta`. Where correction load accrues more slowly than stored order decays, raw `Q` rises without bound in `r` and every incentive attached to raw `Q` is an incentive to over-frame. That is not a moral claim about the observer; it is the arithmetic of what they are being paid in. **Measure `Q^V`, or measure `Z` alongside `Q`, or the ledger you are optimising will point away from contact.** Chapter 40 §40.2 gives the instrument; Chapter 48 shows the same pathology at institutional scale.

35.21 Drift, and why calibration is a maintenance item

The correction map as written in §35.14 has no source term. Left alone, `ln r` sits where it is. But `E` moves — the building stock changes, the codebase changes, the population changes — and a frame that matched the old environment does not match the new one. **In a moving environment, holding your frame constant is drift**, and drift is a forcing function on the correction loop.

Add it:

```
      ln  r_{t+1}  =  (1  −  a)  .  ln  r_t  +  w_t                [POSTULATE —
the driven correction map]

    w_t = distortion injected per cycle by environmental change, in nats
```

This is now exactly the driven first-order system of Chapter 2 — a leaky integrator with a source — and it has the same steady state. Setting `ln r_{t+1} = ln r_t = ln r_inf` with constant `w`:

```
    ln r_inf = w / a                        THE RESIDUAL DISTORTION
```

Residual distortion is the drift rate divided by the correction gain. It is not zero and it cannot be made zero. A maintained observer in a moving environment does not sit at contact; they sit at a fixed offset from it, and the size of the offset is set by two numbers, only one of which they control.

Computed:

drift w (nats/cycle)	gain a = 0.05	0.10	0.20	0.35	0.50
0.02	0.400 MARKED	0.200 MILD	**0.100 CAL.**	0.057 CAL.	0.040 CAL.
0.05	1.000 SEVERE	0.500 MARKED	0.250 MILD	0.143 MILD	**0.100 MILD**
0.10	2.000 SEVERE	1.000 SEVERE	0.500 MARKED	0.286 MILD	0.200 MILD
0.20	4.000 SEVERE	2.000 SEVERE	1.000 SEVERE	0.571 MARKED	0.400 MARKED

The requirement to hold the residual inside the calibrated band is `a >= w / 0.10`:

drift w per cycle	minimum gain required
0.02	0.20 — achievable
0.05	0.50 — achievable, at the top of the working band
0.10	1.00 — deadbeat, forbidden by S35.10 at any material `Q`
0.20	2.00 — at the stability limit; **not achievable at all**

If the environment injects more than about a tenth of a nat of distortion per correction cycle, no achievable gain holds calibration, and the only remaining lever is to shorten the cycle.

That is a hard conclusion and it is the reason the audit cadence is a specification item rather than a preference. You cannot correct your way out of a drift problem by correcting harder; past `a ~ 0.5` the increments buy little, and past `a = 1` they are forbidden by the admissibility condition. **You correct your way out of it by correcting more often.**

The drift rate `w` is not something this manual can give you a number for; it is a property of your domain and it must be estimated locally — by re-auditing a *frozen* frame against a fresh environment and reading the distortion that accumulated between the two dates. What the manual can give you is the cadence rule that follows from the table:

> **CADENCE RULE.** The audit interval **shall** be short enough that the distortion injected between audits is below the tolerance band width. In symbols, with `w_year` the annual drift and `B` the band width in nats, the interval in years satisfies `T <= B / w_year` **and** the gain satisfies `a >= w_year . T / B`. Both, not either.

If your band is 0.10 nats and your domain moves 0.10 nats a year, you need an annual cycle at gain 0.5 or better, or a quarterly cycle at gain 0.13. **Most organisations in fast-moving domains are running frames last validated at a cadence set by convenience, and the residual distortion that implies has never been computed.** It takes four numbers and one line of arithmetic to compute it, and the numbers are `w`, `a`, `B` and the interval.

35.22 Why the two directions are policed unequally

A closing observation about the environment in which all of this has to be done, offered as the author's reading and badged as such.

> **[AUTHOR'S OBSERVATION — not derived.]** Institutions detect and punish under-framing readily and detect over-framing almost never. The reason is visible in the signature panels: **the under-framed observer is slow, uncertain and produces thin output, and every one of those is directly measurable by instruments organisations already run.** The over-framed observer is fast, confident and produces abundant output, and detecting them requires an instrument — off-model seeded items — that almost nobody runs.

The arithmetic supports the asymmetry rather than the intuition. In the worked example the over-framed veteran produces a worse total error count than the untrained novice, in a fifth of the time, while being the office's most productive inspector by the office's own metric. **The state that gets caught is the one that costs least; the state that costs most is the one that gets promoted.** That is not a claim about human nature. It is a claim about which quantities happen to be easy to count, and it is correctable by counting a different one — which is the entire reason S35.4 exists and the entire reason it is the requirement to defend when someone proposes to simplify the audit.

The Specification

S35.1 — Report in log units, signed and unsigned

Distortion **shall** be reported as `ln r` (signed) and `|ln r|` (magnitude), both in nats. **A raw ratio or a percentage shall not be used as a distortion measure**, because it is not symmetric between the two directions and does not compose additively.

S35.2 — Never the arithmetic mean of r

Aggregation across occasions **shall** use mean `|ln r|` for distortion and mean `ln r` for bias. **The arithmetic mean of r shall not appear in any report governed by this manual.**

S35.3 — Two estimators minimum

Any `r` used for a decision **shall** rest on at least two estimators drawn from different rows of §35.13, or on one seeded audit. **A single behavioral estimator shall be recorded as an indication, not a measurement.**

S35.4 — Off-model content at 20 per cent minimum

Restated from S34.9 because it is the requirement most often dropped: **every calibration audit shall contain at least 20 per cent off-model items. An audit without them cannot detect over-framing.**

S35.5 — Separate intake and assertion records

Any observer whose calibration is to be measured **shall** keep an intake record distinct from the assertion record. **Where the two are the same document, `r` is 1 by construction and cannot be measured.** This is a records-design requirement and it must be in place before the audit, not after.

S35.6 — Direction before correction

The sign of `ln r` **shall** be established before any correction is applied. **Δ alone is not sufficient to choose a procedure**, and the two directions have opposite treatments (§35.15 against §35.16).

S35.7 — The reduced-sample test before any training prescription

Where `r < 0.905`, the reduced-sample test **shall** be run before training is prescribed. **An overloaded observer given training gets worse**, because training adds to the load that is already the binding constraint.

S35.8 — The environment measured first

Restated from S34.10: where performance is poor, `E` **shall** be measured independently before any correction is applied to the observer.

S35.9 — Gain inside the stability band

Correction regimes **shall** target `0 < a < 2` and **shall** be reviewed for divergent behavior — alternating sign with growing magnitude — at least quarterly. **A regime showing that signature shall be stopped rather than tuned.**

S35.10 — No deadbeat correction of a large Q

Where stored order is large, correction gain **shall not** be set at or near 1. Verify against the admissibility condition `C_t <= (1−δ)Q_t + ηM_t` before setting the gain.

S35.11 — Per-stage measurement in pipelines

Distortion in a multi-stage process **shall** be measured per stage. Reports **shall** give both net `|ln r_total|` and injected `Σ|ln r_i|`. **A system whose injected distortion materially exceeds its net distortion shall be recorded as fragile** and shall not be certified on its end-to-end figure alone.

S35.12 — Declare the budget

Any process with a calibration requirement **shall** state its total distortion budget in nats and its per-stage allocation, before measurement.

S35.13 — Operating point declared as policy

Any deliberate departure from `r = 1` **shall** be recorded in writing with the loss function that justifies it. **An undeclared departure shall be treated as instrument error, not as policy.**

S35.14 — Audit cadence tied to environmental drift

The audit interval **shall** be shorter than the time in which the environment is expected to move by the width of the tolerance band. The estimate of that rate **shall** be written down and revised.

S35.15 — Consent and decline

Calibration measurement is measurement of a person, and CANON §5.4 governs. Participation **shall** be voluntary, refusal **shall** be recorded as refusal rather than as a result, and results **shall not** be used for a purpose not disclosed before the audit.

The Failure Modes

FM-35.1 — Confidence Audited as Calibration

Description. An observer's certainty, fluency and speed are taken as evidence of contact.

Damages. Detection of over-framing entirely.

Diagnostic signature. Falling surprise rate (E2) with rising false-discovery rate. No seeded audit on file.

Fix. S34.8, S35.4. The symptom of the disease cannot be the test for it.

FM-35.2 — The Model-Shaped Audit

Description. The calibration audit is built from the same checklist the observer works from. Every seeded item is on-model.

Damages. The audit itself. It will certify an over-framed observer as excellent, with high apparent sensitivity.

Diagnostic signature. Compare the seeded item list against the standing model. If the overlap is total, the audit is a mirror.

Fix. S35.4. Off-model items at 20 per cent minimum, generated by someone who does not work from that model.

FM-35.3 — The Arithmetic Mean of Ratios

Description. An observer's `r` values across six occasions are averaged arithmetically and reported as their calibration.

Damages. The measurement. By Jensen, the figure is biased toward over-framing, and an observer who swings both ways is reported as a mild over-framer with no indication of instability.

Diagnostic signature. Any report containing a mean `r`.

Fix. S35.2. Report mean `|ln r|` and mean `ln r`.

FM-35.4 — Compensating Distortion

Description. Two stages with opposite distortions, netting to `r ~ 1`. End-to-end measurement certifies the system.

Damages. Robustness. The system fails without warning on any personnel change.

Diagnostic signature. Per-stage `r` values spread widely around a net near 1; injected distortion much larger than net.

Fix. S35.11. Measure per stage. Correct each stage independently, and expect the end-to-end figure to get *worse* before it gets better — **correcting the over-framed stage first will reveal the under-framed stage that was hiding behind it, and a system that has been in compensation for years will look like it is deteriorating during the repair.** Say this in advance or the repair will be reversed.

FM-35.5 — The Single Measurement Point

Description. A four-stage pipeline is measured at one stage and the result reported as the pipeline's calibration.

Damages. Everything downstream of the measurement point.

Diagnostic signature. Ask where the measurement was taken. In the §35.19 example, a measurement at the write-up returns `ln r = −0.0228` for a pipeline that delivers `r = 2.11`.

Fix. S35.11.

FM-35.6 — Training Prescribed for Overload

Description. An observer with $r < 1$ is sent to training. Their frame was already adequate; the intake rate was the constraint. Training adds load.

Damages. r falls further; C rises; the observer is now also being told the problem is their competence.

Diagnostic signature. The reduced-sample test. If performance improves markedly on a third of the material, this is load.

Fix. S35.7, then Lever D (§35.16): cut the intake rate.

FM-35.7 — The Deframing Prescription

Description. An under-framed observer is advised to have fewer expectations, be more open, stop overthinking. D is reduced further.

Damages. D, therefore M. This is FM-34.5 and it is iatrogenic.

Diagnostic signature. Advice on file that targets the numerator when r is already below 1.

Fix. Lever C. Supply structure: exemplars, a taxonomy, a checklist **with a decision rule**, a forced written conclusion per case.

FM-35.8 — The Checklist That Raises P

Description. An under-framed observer is given a checklist of things to look at. It contains no rule for what any observation implies.

Damages. Raises intake without raising assertion. r falls. The observer is now noticing more and concluding no more.

Diagnostic signature. The checklist's items are all nouns and no item has a "therefore".

Fix. Add the decision rule to each item. **A list of what to look at is an intake instrument; a list of what things mean is a frame instrument; only the second treats under-framing.**

FM-35.9 — Throughput Incentive Read as Skill

Description. Time per case is measured and rewarded. r rises across the whole body over years. The organisation reads it as growing expertise.

Damages. P organisation-wide, therefore r, therefore V.

Diagnostic signature. E4 falling across the whole population simultaneously; off-model detection falling in step; no change in the environment. **A population-wide movement in r is an organisational signal, not a personnel one.**

Fix. Chapter 48. Individual correction will not hold against a standing incentive. Change what is measured, or accept the operating point and declare it under S35.13.

FM-35.10 — Uncorrected r Defended as Judgement

Description. A persistent departure from `r = 1` is described after the fact as professional judgement, prudence, or decisiveness. No loss function was ever stated.

Damages. Correction is blocked, and the state is now protected by an account of itself.

Diagnostic signature. Ask for the loss function. A policy has one, in numbers, from before. An uncorrected instrument produces one on request, in words, afterwards.

Fix. S35.13. Require the declaration in advance. **A stated policy is auditable; a retrospective justification is not, and the difference between them is the whole of the governance available here.**

FM-35.11 — Calibration Assumed to Persist

Description. An observer certified as calibrated four years ago is treated as calibrated now. The environment has moved.

Damages. `r`, silently and progressively.

Diagnostic signature. Date of last audit relative to the estimated environmental drift rate.

Fix. S35.14. **There is no decay term on distortion: nothing returns an observer to contact except correction, so calibration is a maintenance item with a schedule, like any other.**

FM-35.12 — The Divergent Corrector

Description. Effective gain at or above 2. Position reverses fully at each contradiction, then reverses again.

Damages. `r`, with growing amplitude.

Diagnostic signature. Alternating sign of `ln r` across successive audits with increasing magnitude.

Fix. S35.9. Reduce the gain. Check specifically for the same discrepancy being acted on more than once, which is the usual mechanism by which a modest per-event gain compounds past 2.

Worked Example

A building inspector, measured three times across fourteen years

A municipal inspection office wants to know what happens to an inspector's calibration over a career. It has never measured this. It runs one audit design at three career points, using inspectors who volunteered under S35.15.

The audit. A vacant municipal building is prepared with **50 conditions**, adjudicated in advance by two structural engineers and a code officer who take no part in the inspections.

- **35 on-model**: conditions named in the office's standing inspection schedule — egress width, handrail height, panel labelling, the usual.

- **15 off-model**: conditions outside the schedule — a roof drain discharging onto a fire escape landing, a mislabelled shut-off, an added partition that changed the occupancy classification, a masonry repair executed in a mortar harder than the surrounding brick, and eleven others.

Off-model content is 30 per cent, above the S35.4 minimum. **The office's first instinct was to build the audit from the inspection schedule, which would have made every item on-model and guaranteed the result the office expected.** The 30 per cent is the whole reason this audit found anything.

Each inspector carries two records: a **field notebook** (intake) and a **filed report** (assertion). S35.5 is satisfied by existing practice; the office already required both.

Point 1 — Week 2. A new inspector, nine days out of orientation.

```
    conditions present                        50   (35 on-model, 15 off-model)
    real conditions appearing in the notebook 34
    items in the filed report                   12   (9 real, 3 not real)
    time on site                                232 minutes

    P_hat = 34/50 = 0.680000
    D_hat = 12/50 = 0.240000
   r    = 0.352941       ln r = -1.041454     Δ = 1.0415 nats    SEVERE UNDER-
FRAMING
```

Read the numbers together. **The notebook contains 34 of the 50 real conditions — this inspector's intake is not the problem, and it is better than the veteran's will be.** The report contains twelve items. Sixteen real conditions were written down in the notebook and never made it into a finding. Nearly four hours on site.

The subjective report, recorded at the time: *"I saw a lot of things I wasn't sure counted."* That sentence is `r = 0.35` in plain language, and it is exactly Chapter 34's Limit 2 — intake without frame, the moment collapsing to a third of the environment.

Veridical throughput: `V/E = min(r,1) = 0.3529`. **Sixty-five per cent of the building was missed, and almost all of it was missed after being seen.**

Point 2 — Month 9. The same inspector, after training and supervised fieldwork.

```
    real conditions appearing in the notebook 42
    items in the filed report                   44   (40 real, 4 not real)
    time on site                                118 minutes

    P_hat = 42/50 = 0.840000
    D_hat = 44/50 = 0.880000
   r    = 1.047619       ln r = +0.046520     Δ = 0.0465 nats    CALIBRATED
```

Both terms rose. **This is the case Chapter 34 §34.5 flagged: D and P went up together, which changed what the inspector could handle without changing how distorted they were — except that the frame rose slightly faster, which moved r from 0.353 to 1.048.** Off-model detection is 60 per cent. False discovery rate 9 per cent. True conditions found: 40 of 50, up from 19. Half the time on site.

This is what training is for and it worked. Note what it treated: **the correction raised D, which was the correct lever because the diagnosis was under-framing with adequate intake, and the reduced-sample test would have confirmed it.** Had the diagnosis been overload, this same intervention would have made things worse.

Point 3 — Year 14. The same inspector, now the office's most productive.

```
    real conditions appearing in the notebook 30
    items in the filed report                 57   (30 real, 27 not real)
    time on site                              41 minutes

    P_hat = 30/50 = 0.600000
    D_hat = 57/50 = 1.140000
  r   = 1.900000     ln r = +0.641854    Δ = 0.6419 nats    MARKED, edge of SEVERE
```

The intake fell below where it was in week two. Thirty of fifty conditions reach the notebook, against thirty-four in the first week. The report has fifty-seven items, of which twenty-seven are not there.

The signature panel:

Metric	Week 2	Month 9	Year 14
r	0.353	1.048	1.900
ln r	-1.0415	+0.0465	+0.6419
Band	Severe under	**Calibrated**	Marked over
Intake P_hat	0.680	0.840	**0.600**
On-model hit rate	0.400	0.886	**0.943**
Off-model hit rate	0.333	0.600	**0.133**
False discovery rate	0.250	0.091	**0.474**
Minutes on site	232	118	**41**
True conditions found	19/50	40/50	35/50
Veridical throughput V/E	0.353	1.000	1.000
Confabulated fraction	—	—	**0.474**

Two rows carry the diagnosis. **On-model detection is at its career maximum, 94 per cent — and off-model detection has fallen to 13 per cent, below the novice's 33 per cent.** The inspector has become excellent at finding what the schedule contains and worse than a nine-day novice at finding what it does not. That is FM-34.4, and it is the expertise-decay subtype of §35.8: a frame that was correct, earned honestly, and outlived the building stock it was built on.

Productivity, priced properly. The office's metric is inspections per day, and by that metric Year 14 is 5.7 times Week 2.

	true conditions found per hour	false alarms per hour
Week 2	2.328	0.776
Month 9	20.339	2.034
Year 14	**43.902**	**39.512**

The veteran finds real conditions nineteen times faster than the novice and generates false alarms fifty-one times faster. Under equal error costs, the veteran's total error count on this building — 20 misses plus 27 false alarms, 47 — is worse than the novice's 44, achieved in a fifth of the time. The trained mid-career point, at 14, is a third of either.

Costed:

Operating point	r	misses	false alarms	cost at 1:1	at 10:1 (miss:alarm)	at 1:3 (miss:alarm)
Week 2	0.353	41	3	44	413	50
Month 9	1.048	10	4	**14**	**104**	**22**
Year 14	1.900	20	27	47	227	101

The calibrated point wins under every cost structure tested. **Note that the office would have to value a miss at more than about six times a false alarm before the veteran's operating point beats the novice's, and no cost structure makes either beat the calibrated one.**

Storage consequence. With office-typical personal parameters $\eta = 0.35$, $\delta = 0.08$, $C = 0.20$, and E normalized to 1.0 nat:

Point	M = rE	Q_inf on M	Q_inf on V	stored error
Week 2	0.353	**-0.956**	-0.956	0.000
Month 9	1.048	+2.083	+1.875	+0.208
Year 14	1.900	**+5.813**	+1.875	**+3.938**

The novice's steady state is negative — $\eta M < C$, the unsustainable regime Paper E Theorem 3.5 declines to hide — which is a correct model statement about a person who cannot continue in that condition and must be moved out of it, in either direction, quickly.

The veteran's steady state is the largest of the three and **68 per cent of it is stored error.** Fourteen years of confident, fast, internally consistent structure, two thirds of which the building did not supply. It will be defended, and the defence will be sincere, because from inside there is no difference between the earned third and the manufactured two thirds.

The correction. Diagnosis: over-framed, expertise-decay subtype, driven by a throughput incentive (FM-35.9 present at the office level — E4 has fallen across the whole inspection staff).

Because the driver is organisational, S35.13 applies first: the office must either change the throughput metric or declare the operating point as policy with a loss function. It changes the metric, adding off-model detection rate to the quarterly report.

Individual procedure, both levers:

- **Lever B (raise P), primary.** A mandated fifteen-minute off-schedule pass with its own notebook section, question fixed as: *what is in this building that the schedule does not cover*. Rotation onto two building types outside the inspector's usual stock, one per quarter.
- **Lever A (lower D), secondary.** Every filed item must cite the notebook entry that generated it. Items with no notebook entry cannot be filed.

Target gain `a = 0.35` — the working band, half-life 1.61 cycles, comfortably inside the stability condition and nowhere near deadbeat, which S35.10 forbids at this `Q`.

Projected trajectory from `ln r = +0.6419`:

cycle	ln r	r	band
0	+0.6419	1.900	Marked
1	+0.4172	1.518	Marked
2	+0.2712	1.312	Mild
3	+0.1763	1.193	Mild
4	+0.1146	1.121	Mild
5	+0.0745	1.077	**Calibrated**
6	+0.0484	1.050	Calibrated

Five audit cycles to the calibrated band. At the office's quarterly cadence that is fifteen months, and the office should expect the inspector's headline productivity to fall by roughly half during it — time on site rising from 41 minutes toward something near the mid-career 118 — while true conditions found per inspection rises and false alarms collapse.

Say that in advance. A correction that reduces the number people have been watching, in order to improve numbers nobody has been watching, is reversed in month four unless the trade was written down in month zero. That is not a psychological observation; it is the reason S35.12 requires the budget to be declared before measurement rather than after.

Falsification (FVP-6). The office preregisters: the inspector's frozen post-correction `r` predicts held-out false-discovery rate below 0.15 and off-model hit rate above 0.40 on a fresh building, and beats (i) persistence from the Year 14 audit, (ii) a model using building complexity alone, (iii) an autoregressive model on the inspector's own past filed-item counts, and (iv) **the same model fitted to inspector/building pairings shuffled at random across the office's staff.**

If the shuffled control predicts as well, `r` is a property of the building and not of the inspector, the instrument measures nothing about the person, and the office has agreed in writing to publish that internally and stop using it.

That commitment is the difference between an audit and an opinion, and this office now has one on file for the first time in its history.

Chapter 34 gave the moment its equation and Chapter 35 gave its distortion a measure. Both rest on a choice that has been asserted twice and justified only in passing: that the measure is a logarithm. Chapter 36 justifies it, and in doing so answers the question this volume was written around — what mathematical functions actually mean when they are read through the lens perception uses.

Chapter 36 . How Mathematical Functions Operate From a Perceptual Lens

Two chapters have now asserted that distortion is measured by a logarithm, and neither has fully earned it. Chapter 35 gave three reasons and promised this chapter would give the deep one.

Here it is. **Perception does not report the world's quantities. It reports the world's ratios, and a device that reports ratios as though they were differences is a logarithm.** Everything in this chapter follows from that one sentence, and what follows from it is not a philosophical position. It is a complete translation table between the arithmetic you do on paper and the arithmetic the recipient of your design is actually running.

The consequence that organises the chapter, stated at the front so that nothing after it is a surprise:

```
ADDITION IN PERCEPTION IS MULTIPLICATION IN THE WORLD.

perceived +1 unit   ->   world x10
perceived +2 units  ->   world x100

LINEAR IN THE LENS IS EXPONENTIAL IN REALITY.
```

A designer who promises a linear improvement in a quantity the recipient perceives logarithmically has promised an exponential and does not know it. That failure has a name and a price and both are computed in §36.22.

The Law

36.1 The Lens Theorem

> **THE LENS THEOREM [CONDITIONAL THEOREM].** *If* the map from a world quantity `I` to its perceived magnitude `S` is continuous and converts multiplication into addition — that is, if
>
> ```
> S(I_1 . I_2) = S(I_1) + S(I_2)              for all I_1, I_2 > 0
> ```
>
> *then* that map is a logarithm, and it is the only such map up to a single scale constant:
>
> ```
> S(I) = k . ln I                              k > 0
> ```
>
> Equivalently, with a reference intensity `I_0` at which the percept is zero,
>
> ```
> S(I) = k . ln(I / I_0)
> ```

This half is a theorem and it is not new: it is the continuous solution of Cauchy's logarithmic functional equation, and it is exactly the statement that the logarithm is the unique continuous isomorphism from the multiplicative group of positive reals to the additive group of reals. No psychology is involved in proving it. It is an algebraic fact about the two operations.

The theorem is *conditional*, and the antecedent — that perception is such a map — is an empirical claim about human observers. §36.3 gives the century of evidence for it and §36.5 gives the honest limits. **The badge is on the antecedent, not on the inference.** This matters, because it means the argument of this chapter cannot be defeated by disputing the mathematics; it can

only be defeated by defeating Weber and Fechner, and a great deal of the chapter survives even then (§36.6).

The inverse direction is the one designers need:

```
I = I_0 . e^{S/k}
```

A linear schedule of percepts requires an exponential schedule of world quantities. If you want someone to experience equal steps, you must deliver equal ratios, and equal ratios compound. That is the whole practical content of the theorem and everything in The Engineering is a worked consequence of it.

36.2 Weber's law, and the derivation of Fechner's

Weber (1834): the just-noticeable difference in a stimulus is a constant fraction of the stimulus.

```
dI / I = c              c = the Weber fraction, constant over a working range
```

Fechner (1860) took the step that makes it a lens. Postulate that every just-noticeable difference is one unit of subjective magnitude — that JNDs are perceptually equal — and integrate:

```
dS = k . dI / I

S = k . int dI/I = k . ln(I) + const

S = k . ln(I / I_0)                    with I_0 the threshold where S = 0
```

Fechner's postulate — that all JNDs are subjectively equal — is a postulate, and this manual badges it wherever it is used. It is not derivable from Weber's law. It is an additional assumption, it has been disputed since it was made, and Fechner's law is what you get if you accept it.

Verified numerically, `k = 1`, `I_0 = 1`, integrating to `I = 100` by 200 000 midpoint rectangles:

```
midpoint Riemann sum = 4.605170176
ln(100)              = 4.605170186
absolute error       = 1.021e-08
```

The Weber fraction determines how many perceptual steps a range contains, and the numbers are sobering:

Weber fraction c	JNDs across a 2:1 range	JNDs across a 1000:1 range
0.02	35.00	348.83
0.05	14.21	141.58
0.10	7.27	72.48
0.20	3.80	37.89

A thousandfold range of world quantity contains 72 discriminable steps at a Weber fraction of 0.10, and a doubling contains seven. The world's dynamic range is enormous and the lens's is small; that compression is the lens's whole job, and complaining that people cannot tell a 3 per cent improvement from nothing is complaining that they have working sense organs.

36.3 Stevens' power law, and why it does not defeat the argument

Stevens (1957) proposed a competitor:

```
psi = k . I^a                              the power law
```

and produced it by **magnitude estimation** — asking subjects to assign numbers proportional to their sensations — which yields ratio-scale measurements of subjective quantity and replicates. Chapter 38 treats this in full; it is the existence proof on which the whole of Volume Two rests.

The competition between Fechner and Stevens is real and unresolved. **It does not affect the Lens Theorem's usefulness, and the reason is one line of algebra:**

```
psi = k . I^a        =>        ln psi = ln k + a . ln I
```

A power law is a straight line in log-log coordinates. It is the logarithmic lens with a gain. Under Fechner the exponent is fixed at the boundary case a = 1 in log coordinates; under Stevens it is a measured constant that varies by modality. Either way the operative coordinate is the logarithm and every operation in §36.8's table holds, with the world-side ratio raised to the power a:

```
perceived ratio = (world ratio)^a
perceived change in nats = a . ln(world ratio)
```

The exponent a is a measured conversion key between the physical and perceptual registers (CANON_INTERNAL §3), and that is the honest way to hold both theories at once: the lens is logarithmic, and a is its gain.

Published, replicated exponents, presented as **nominal, modality- and method-dependent** — they are constants of a procedure, not constants of nature:

Modality	exponent a	character
Loudness (1 kHz, sound pressure)	~0.67	compressive
Brightness	~0.33	compressive
Visual length	~1.0	linear
Taste (salt)	~1.3	expansive
Heaviness (lifted)	~1.45	expansive
Warmth on arm	~1.6	expansive
Electric shock	~3.5	strongly expansive

a < 1 compressive, a = 1 linear, a > 1 expansive. **Most senses compress; pain expands.** §36.26 shows that the expansive modalities invert every design rule in this chapter, which is why the exponent must be looked up rather than assumed.

36.4 FVP table — Fechner's law

```
FVP TABLE — FECHNER'S LAW          S = k ln(I / I_0)
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units; ratios dimensionless.	`I/I_0` dimensionless as required — the argument of a logarithm must be. `[S] = [k]`, and k carries whatever unit the percept scale is declared in. **A logarithm of a dimensioned quantity is a category error and appears surprisingly often in print.**	**PASS**
FVP-2 SCALE TYPE	What scale is S on?	`I` is ratio-scaled (physical intensity, true zero). `S` is **interval-scaled**: differences in S are meaningful, but S has no true zero because `I_0` is a convention. **Consequence: you may subtract percepts and you may not divide them.** This is a real and frequently violated constraint.	**PASS**, with the division prohibition
FVP-3 LIMITING CASES	`I -> 0`, `I -> infinity`.	`I -> 0` gives `S -> -infinity`: the lens has no floor. Real: this is why an absolute threshold `I_0` must exist, §36.18. `I -> infinity` gives `S -> infinity` but slowly; real, and it is why saturation and damage limits bound the useful range rather than the lens doing it.	**PASS**, with the floor caveat

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	Invariant.	`S = 0` at `I = I_0`. **Zero percept is the reference intensity, not the absence of the stimulus.** The most consequential fixed point in the chapter and §36.13 is about it.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable parameters?	k and `I_0` are **not separately identifiable from S alone**: `k ln(I/I_0) = k ln I − k ln I_0`, so a change in `I_0` is absorbed by an additive offset. Only k and the offset are recoverable. Fixing `I_0` by an independent threshold measurement resolves it.	**PASS**, conditional on independent threshold measurement
FVP-6 FALSIFICATION	Held-out prediction against baselines.	**Extensively tested for over 160 years, and the result is mixed:** Fechner's law fits many modalities well over the middle of their range and is beaten by Stevens' power law in others, particularly at range extremes. It is falsifiable, it has been tested, and it has partly failed.	**RUN — PARTIAL PASS. Honest verdict: a good approximation over working ranges, not a universal law.**

Gate	Test applied	Finding	Verdict
FVP-7 REPRODUCIBILITY	Cross-individual, cross-lab recovery.	Weber fractions and Fechnerian fits recover across individuals and laboratories within modality, with individual variation in `c` and `k`. **This is the gate psychophysics actually passes, and it is why psychophysics is the model for this volume.**	**RUN — PASS within modality, with stated individual variation**

Two gates run and passed, one run and partially passed. No formula elsewhere in Volume Two has that record, and that is precisely why this chapter is where the volume's methodological argument lives.

36.5 FVP table — Stevens' power law

`FVP TABLE — STEVENS' POWER LAW  psi = k . I^a`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units.	`I` must be normalized to a reference before exponentiation, exactly as in Fechner. a dimensionless. `[psi] = [k]`.	**PASS**, with the normalization requirement
FVP-2 SCALE TYPE	Does magnitude estimation give a ratio scale?	**Yes — that is the method's central claim and the reason it matters here.** Subjects assign numbers proportional to sensation; the resulting scale supports multiplication and division. This is the existence proof that subjective quantity can be ratio-scaled.	**PASS** — and this row is the single most important row in Volume Two

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	`I -> 0`, `I -> infinity`.	`I -> 0` gives `psi -> 0` for `a > 0`: unlike Fechner, the power law has a floor at zero. `I -> infinity` diverges, bounded in practice by damage limits.	**PASS**
FVP-4 FIXED POINTS	Invariant.	`psi = k` at `I = 1` (the reference). In log-log the invariant is the intercept.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	a recovers as the slope of a log-log regression and is robust. k depends on the numerical instruction given to subjects and **is not a property of the observer.**	**PASS for a; k is procedure-dependent**
FVP-6 FALSIFICATION	Tested?	Extensively. Exponents replicate within method. **They vary with method, stimulus range, and the instruction given** — a documented and acknowledged limitation, not a hidden one.	**RUN — PASS with declared method-dependence**
FVP-7 REPRODUCIBILITY	Cross-lab.	Exponents recover across laboratories within a fixed procedure. **The honest statement: report a Stevens exponent as a measured constant of a procedure, not as a constant of nature.**	**RUN — PASS, procedure-relative**

36.6 FVP table — the Lens Theorem itself

```
FVP TABLE — THE LENS THEOREM      S(I_1 I_2) = S(I_1) + S(I_2)  =>  S = k ln I
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units.	The functional equation is dimensionless in its argument; `I` must be a ratio to a reference throughout.	**PASS**
FVP-2 SCALE TYPE	What is required of the inputs?	The world quantity must be **ratio-scaled** for multiplication of stimuli to be defined at all. Where the world quantity is only ordinal, the theorem does not apply and no log should be taken.	**PASS**, conditional
FVP-3 LIMITING CASES	Degenerate cases.	`k -> 0` collapses all percepts to zero — no discrimination. `k -> infinity` makes every difference infinite — no constancy. Both are real failure modes of sensory systems and both are named.	**PASS**
FVP-4 FIXED POINTS	Invariant.	`I = 1` (the reference), where `S = 0`. **Multiplication's identity maps to addition's identity. The two neutral elements correspond, and that is what makes it an isomorphism rather than a mere resemblance.**	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	k recovers from any two (I, S) pairs. Uniqueness of the log form is proved, not fitted.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	What would kill it?	The *inference* cannot be falsified — it is algebra. **The antecedent can:** demonstrate a modality in which equal perceptual steps correspond to equal world differences rather than equal world ratios, robustly and across the range, and the lens does not apply to that modality. Visual length, at a $\sim$ `1.0`, is close to being such a case.	**The conditional: PASS by proof. The antecedent: falsifiable, tested, holds in most modalities and not all.**
FVP-7 REPRODUCIBILITY	Cross-individual.	Inherits from Fechner and Stevens: the *form* recovers robustly; the constants c, k, a vary between individuals and must be measured per person where precision is needed.	**RUN — form reproduces, constants are individual**

36.7 The honesty clauses for this chapter

1. **The Lens Theorem's mathematics is a theorem; its applicability is an empirical claim.** Do not present the second with the confidence of the first.
2. **Fechner's equal-JND postulate is a postulate.** It does not follow from Weber's law.
3. **Stevens' exponents are constants of a procedure.** Report the method, the range and the instruction alongside any exponent, or the number is not usable.
4. **Percepts are interval-scaled, not ratio-scaled, under Fechner.** You may subtract them. You may not divide them. Under Stevens' magnitude estimation you may do both, and that is the methodological difference between the two that matters most to this manual.
5. **Nats of one modality are not nats of another.** S36.9 is explicit about this and it is the most likely way for a reader to misuse this chapter.
6. **This chapter describes magnitude judgment.** It is not a theory of cognition, of preference, or of value, and it should not be extended to those without separate evidence.

The Engineering

36.8 The translation table

This is the section the chapter exists for. Every arithmetic operation, read on both sides of the lens.

```
LENS SIDE  = what the observer experiences, S = k ln(I/I_0)
WORLD SIDE = what you actually have to build, I
```

Operation in the LENS	Corresponding operation in the WORLD
`S_a + S_b` — addition	**Multiplication.** `I_a . I_b / I_0`
`S_a - S_b` — subtraction	**Division.** The ratio `I_a / I_b`
`0` — the origin	**The reference `I_0`. Not absence.**
a negative percept	`I < I_0`: a reciprocal ratio, below reference
`c . S` — scalar multiplication	**Exponentiation.** `I^c` (relative to reference)
`S / c` — scalar division	**The c-th root.** `I^(1/c)`
`2S` — doubling the percept	**Squaring the world quantity**
arithmetic mean of percepts	**Geometric mean** of world quantities
midpoint between two percepts	**Geometric midpoint**, `sqrt(I_a . I_b)`
linear interpolation	**Geometric interpolation**
arithmetic progression (equal steps)	**Geometric progression (equal ratios)**
sum over n terms	**Product over n terms**
`|S_a - S_b|` — distance	**`|ln(I_a/I_b)|` — the log-ratio metric**
variance / normal distribution	**Log-variance / lognormal distribution**
`dS/dt` — a rate of change	**A proportional growth rate**, `(dI/dt)/I`
`dS/dx` — a derivative	**An elasticity**, `(dI/I)/(dx/x)` when x is also logged
a straight line, rising	**Exponential growth**
a straight line in log-log	**A power law**
a concave curve in the world's own units	**A straight line, seen through the lens**
the integral `intS dt`	**The geometric (product) integral of I**
unbounded below, `S -> -infinity`	**`I -> 0`: the missing floor, §36.18**

Twenty rows. Each of the sections that follow takes the ones that carry weight.

36.9 Addition: the founding row

```
S(I_a) + S(I_b) = k ln(I_a/I_0) + k ln(I_b/I_0) = k ln( I_a I_b / I_0^2 )
```

Two percepts added correspond to the two world quantities multiplied, with the reference divided out once. The clean form, in the base-10 lens people actually use when they say "orders of magnitude":

perceptual step	world multiple
+0	x1
+1	x10

perceptual step	world multiple
+2	x100
+3	x1 000
+4	x10 000
+5	x100 000
+6	x1 000 000

In nats:

step (nats)	world multiple
+0.250000	x1.284025
+0.500000	x1.648721
+0.693147	**x2.000000**
+1.000000	x2.718282
+2.000000	x7.389056
+3.000000	x20.085537
+4.605170	x100.000000

One octave — a halving or a doubling — is 0.693147 nats, which is 0.301030 decades. That single constant, `ln 2`, is the conversion between the language of ratios and the language of steps, and it appears in this manual as the marked/severe distortion boundary (Chapter 35), as the pitch octave (CANON_SONIC), and as the surplus gained by doubling `N` (Chapter 8). **It is the same constant doing the same job in four places, because there is only one lens.**

36.10 The mean: an average of percepts is a geometric mean of quantities

The single most common arithmetic error made about perceived quantities.

```
mean of percepts:   (S_a + S_b)/2 = k ln( sqrt(I_a . I_b) / I_0 )
```

The world quantity corresponding to the average percept is the geometric mean, not the arithmetic mean.

between	arithmetic midpoint	**perceptual midpoint**
1 and 100	50.5000	**10.0000**
10 and 1 000	505.0000	**100.0000**
1 and 2	1.5000	**1.4142**
300 and 2 400 lux	1 350.0000	**848.5281**
20 Hz and 20 000 Hz	10 010.0000	**632.4555**

Read row one and row five. **The quantity that sits halfway between 1 and 100 in experience is 10, not 50.5.** The frequency halfway between the bottom and the top of human hearing is 632 Hz — which is roughly middle of the piano, and any musician would tell you so without the arithmetic, because musicians have been working in the lens for four hundred years and label their instrument in octaves.

The corollary for ratios and rates:

```
+50% then −50%:     1.5 x 0.5 = 0.75            net −25.00%
arithmetic mean of the two changes:              0.00%     WRONG
geometric mean factor: sqrt(0.75) = 0.866025    −13.40% per step   CORRECT
mean of the logs: (ln1.5 + ln0.5)/2 = −0.143841 nats -> exp = 0.866025
```

A gain of 50 per cent followed by a loss of 50 per cent is a loss of 25 per cent, and the average of the two changes is not zero but -13.40 per cent per step. This is not a subtlety. It is the reason Chapter 35's aggregation rule forbids the arithmetic mean of a calibration ratio, and it is the same error there as here.

A four-period example: factors 1.10, 1.10, 0.80, 1.25 give a net of 1.210000, a geometric mean of 1.048809 per period, and an arithmetic mean of factors of 1.062500 — which, compounded over four periods, would predict 1.2744 against the true 1.2100. **The arithmetic mean of ratios always overstates, by Jensen's inequality, and always in the same direction.**

36.11 The metric: distance is a log-ratio

```
perceptual distance between I_a and I_b = |ln(I_a) − ln(I_b)| = |ln(I_a/I_b)|
```

a	b	difference	ratio	perceptual distance (nats)
1	2	1	2.0000	**0.693147**
50	100	50	2.0000	**0.693147**
1 000	2 000	1 000	2.0000	**0.693147**
1	1.1	0.1	1.1000	**0.095310**
100	110	10	1.1000	**0.095310**

Equal perceptual distance, differences varying by a factor of a thousand. The gap from 1 000 to 2 000 is the same experience as the gap from 1 to 2. This is the row that makes every absolute-difference metric wrong for a perceived quantity, and it is why Chapter 35's distortion measure is `|ln r|` and not `|r − 1|`.

36.12 The derivative: a slope in the lens is an elasticity in the world

```
d(ln psi) / d(ln I) = (I / psi) . (dpsi/dI) = elasticity
```

For Stevens' power law `psi = k I^a`, this evaluates to exactly a. **The Stevens exponent is the elasticity of percept with respect to stimulus — the percentage change in experience per percentage change in the world — and it is constant across the whole range.** That is what "the log-log plot is a straight line" means, restated so that an engineer can use it.

For the Moment Equation of Chapter 34, the same operation gave unit elasticities in all three terms: `dM/M = dD/D − dP/P + dE/E`. **That derivation and this one are the same operation, which is the strongest internal evidence that Volume Two's formulas are all in the same register.**

36.13 Zero in the lens is the reference in the world

`S = 0` at `I = I_0`, and `I_0` is not zero. **The origin of the perceptual coordinate is not the absence of the stimulus. It is the reference level.**

The consequences are practical and they are everywhere:

- **A percept of zero means "at reference", not "nothing there".** Silence in a room is not zero sound; it is sound below threshold.
- **Percepts may be subtracted and may not be divided** (FVP-2 above). "Twice as bright" is a statement about the ratio-scaled magnitude estimate, not about the Fechnerian percept, and the two are different objects.
- **The reference is a design decision.** It is set by adaptation, by context, and by what the observer was exposed to before they arrived. **A designer who does not control the reference is not controlling the percept**, and controlling the reference is generally cheaper than moving the quantity — which is why an entry vestibule that is dim makes the reading room beyond it feel bright, and why that is a real engineering technique and not a trick.

36.14 Subtraction: the difference between two readings is the only data

Subtraction in the lens is division in the world, and this row has a name in Book One. Node 08 ECHO — the world's answer is never the voice you sent, and the difference is the only data — is formalized in CANON_INTERNAL §8 as

```
THE RETURN DIFFERENCE      Delta_R = | ln( M' / M ) |
```

where `M` is what was sent and `M'` is what came back. **It is a log-ratio for the same three reasons Chapter 35's distortion measure is: symmetry between the two directions, additivity along a chain, and membership of the same register.** It is also, structurally, the same object as `|ln r|`, which is not a coincidence — both measure how far a returned quantity sits from the one that went out, and the ratio is the only comparison that treats "three times as much" and "a third as much" as equally informative.

sent M	returned M'	ratio	`Delta_R` (nats)
1.0	1.0	1.0000	**0.000000**
1.0	1.2	1.2000	0.182322
1.0	0.8	0.8000	0.223144
1.0	3.0	3.0000	1.098612
1.0	0.333333	0.3333	1.098612

Read the last two rows: **a threefold amplification and a threefold attenuation carry exactly the same amount of information about the world — 1.098612 nats — and an absolute-difference measure would have priced them at 2.0 and 0.667.** Read the first row: **zero return difference means no information came back**, which is not the same as a good result and is often mistaken for one. A system that returns exactly what was put into it has told you nothing except that it is a mirror.

Chapter 43 develops the return difference as an instrument. What belongs here is only the demonstration that it is row two of the translation table and required no separate invention.

36.15 Variance: symmetric in the lens is skewed in the world

If percepts are normally distributed, world quantities are **lognormally** distributed, and the lognormal has a property that catches out everybody who has not met it.

With `ln I` normal about `mu = ln(300)`:

sigma (nats)	median	mean	mode	mean/median
0.2	300.000	306.060	288.237	1.0202
0.5	300.000	339.945	233.640	1.1331
0.8	300.000	413.138	158.188	1.3771
1.2	300.000	616.330	71.078	**2.0544**

The mean, the median and the mode are three different numbers, and the gap between them grows with the spread. At `sigma = 1.2` the mean is more than twice the median and nearly nine times the mode. Reporting "the average" of a perceptually distributed quantity therefore reports a value that almost nobody experiences, and reports it high.

And the interval symmetric in the lens is not symmetric in the world. At `sigma = 0.5` about `mu = ln 300`:

```
−1 sd : 181.959      median : 300.000      +1 sd : 494.616

distance below the median : 118.041
distance above the median : 194.616
```

One standard deviation down is 118 units; one standard deviation up is 195 units. Same perceptual distance, different world distance, by a factor of 1.65. Any error bar, tolerance band or specification range on a perceived quantity that is written symmetrically in world units is asymmetric in experience, and vice versa. **Write the band as a ratio — "x/÷ 1.65" rather than "±195" — and it is symmetric in both.**

This is also why Chapter 35's distortion bands are stated as `|ln r|` intervals and then back-translated to two separate `r` ranges: the band `0.35 <= |ln r| < 0.69` is `r` in [1.419, 1.994) *or* `r` in (0.502, 0.705], and those two intervals have different widths in `r` and identical widths in the lens.

36.16 Linear in the lens is exponential in reality

If the percept is to rise linearly in time,

```
S(t) = S_0 + c.t      =>      I(t) = I_0 . e^{(S_0 + ct)/k}
```

Priced. Suppose the target is a constant perceived step of one Weber-ish increment per period — the world quantity growing by 10 per cent each period, worth `ln 1.1 = 0.095310` nats of percept:

period	world multiple	cumulative percept (nats)
0	1.0000	0.00000
3	1.3310	0.28593

period	world multiple	cumulative percept (nats)
6	1.7716	0.57186
9	2.3579	0.85779
12	**3.1384**	1.14372
24	**9.8497**	2.28744
48	**97.0172**	4.57488

Holding a modest, barely-noticeable rate of perceived improvement for four years of quarterly periods requires the world quantity to grow ninety-seven-fold. The percept has advanced 4.57 nats, which is under two decades of world quantity — and that is the *cheap* interpretation.

Steady perceived improvement requires exponentially growing world quantity, forever, and no physical or budgetary system supplies that. This is the structural reason that every programme of continuous perceived improvement eventually stalls, and it is a fact about the lens rather than about the programme's competence. It is worth saying to the people running such a programme, in those words, because the alternative explanation they will reach for is that they have got worse.

The design response is not to try harder. It is to change the reference (§36.13 and §36.25), or to move to a different modality where headroom remains, or to stop promising a rate that the arithmetic forbids.

36.17 Accumulation: the storage law is a geometric accumulator wearing lens coordinates

The translation table's last structural row is the integral: a sum in the lens is a product in the world. The manual's central recursion is a sum, and it therefore has a world-side reading that has been available since Chapter 1 and has never been written down.

```
Q_{t+1} = (1 − d) Q_t + eta S_t − C_t            the storage law, LENS coordinates
```

Exponentiate both sides. Write `G_t = e^{Q_t}` for stored order in world coordinates, and recall that `S_t = ln N_eff,t` — the surplus is already a logarithm, so `e^{S_t} = N_eff,t` exactly:

```
G_{t+1} = G_t^{(1−d)} . (N_eff,t)^{eta} . e^{−C_t}   the storage law, WORLD coordinates
```

Verified against direct computation, with `delta = 0.08`, `eta = 0.35`, `C = 0.20`, `N = 12`, `u = 0.5` (so `N_eff = 6.5` and `S = 1.871802`), from `Q_0 = 0`:

t	Q_t	G_t = e^{Q_t}	direct e^{Q_(t+1)}	multiplicative form	difference
0	0.00000	1.00000	1.57638	1.57638	0.00e+00
1	0.45513	1.57638	2.39612	2.39612	8.88e-16
2	0.87385	2.39612	3.52216	3.52216	4.44e-16
4	1.61348	5.02024	6.95549	6.95549	8.88e-16
8	2.76936	15.94848	20.14478	20.14478	0.00e+00

Machine precision. Now read the three terms against the translation table:

- **`G_t^(1−d)` — decay is taking a root.** The leak `delta` does not subtract from the world quantity; it takes the `(1−delta)`-th power of it. That is the table's "scalar division in the lens is a root in the world" row, appearing in the manual's most-used equation.
- **`(N_eff)^eta` — conversion is exponentiation.** The efficiency `eta` is not a fraction of an amount added; it is the power to which the effective count is raised before being multiplied in.
- **`e^(−C)` — cost is a divisor.** Dissipation divides the world quantity rather than subtracting from it.

Stored order, in world coordinates, is an effective count, and the storage law is its weighted geometric accumulation. The fixed point follows immediately: setting `G_{t+1} = G_t` gives `G^delta = (N_eff)^eta . e^{−C}`, hence

```
G_inf = ( (N_eff)^eta . e^{−C} )^(1/delta)
```

which for the parameters above evaluates to **295.6376**, identical to `e^{Q_inf}` with `Q_inf = (eta.S − C)/delta = 5.689135`. The two forms agree to fourteen significant figures.

This is not a new theorem — it is Paper E Theorem 3.5 in different coordinates — and it makes no new claim. **What it demonstrates is that the manual has been working in the lens all along.** The storage law looks like a linear accumulator because it is written in log coordinates; the thing it accumulates is a count, and counts combine by multiplication. Every intuition about `Q` being "added to" is an intuition about `G` being multiplied, and that is why Chapter 8's `ln N` ceiling is so brutal, why Chapter 6's headroom shrinks as success accumulates, and why an institution's stored order is very much harder to double than to establish.

36.18 Paper A's surplus law is a Fechner law

This section is the chapter's centre of gravity and it was not put in by hand — it fell out.

Paper A, Theorem 2.1 gives the surplus functional `f(u) = ln(1 + (N−1)u)`. Set `u_0 = 1/(N−1)` and rewrite:

```
f(u) = ln(1 + (N−1)u)
     = ln( (N−1) . ( u + 1/(N−1) ) )
     = ln( (u + u_0) / u_0 )                    with u_0 = 1/(N−1)
```

Compare Fechner:

```
S(I) = k . ln( I / I_0 )
```

and the near-threshold form in which the reference is added rather than divided out, `S = k ln((I + I_0)/I_0)`, which is the standard correction applied when the stimulus approaches the absolute threshold.

They are the same function. The surplus law is Fechner's law with non-parallelism `u` as the stimulus, `k = 1`, and an absolute threshold of `u_0 = 1/(N−1)`.

Verified to machine precision across `N` and `u`:

N	u_0 = 1/(N-1)	u	ln(1+(N−1)u)	ln((u+u_0)/u_0)	difference
2	1.00000000	0.50	0.4054651081081	0.4054651081081	0.00e+00
3	0.50000000	0.20	0.3364722366621	0.3364722366621	0.00e+00
5	0.25000000	1.00	1.6094379124341	1.6094379124340	0.00e+00
12	0.09090909	0.05	0.4382549309310	0.4382549309311	-1.11e-16
50	0.02040816	0.50	3.2386784521643	3.2386784521644	-4.44e-16
1 000	0.00100100	1.00	6.9077552789826	6.9077552789820	0.00e+00

The residuals are floating-point noise. **The identity is exact.**

Three consequences, and they are not decorative.

First, `u_0 = 1/(N−1)` is an absolute threshold of difference. It is the amount of non-parallelism below which the difference does not register as difference. And it *shrinks as N grows*:

N	u_0	f(0.01)	ln N	f(0.01)/ln N
2	1.000e+00	0.009950	0.693147	0.0144
10	1.111e-01	0.086178	2.302585	0.0374
100	1.010e-02	0.688135	4.605170	0.1494
10 000	1.000e-04	4.615022	9.210340	0.5011
1 000 000	1.000e-06	9.210439	13.815511	0.6667

An observer with a larger repertoire has a lower threshold of detectable difference, and the same small non-parallelism `u = 0.01` that is nearly invisible at `N = 2` delivers half the available surplus at `N = 10 000`. That is a precise statement of what expertise buys, and it is the reason the same room, the same conversation and the same manuscript are richer to someone who can tell more kinds of thing apart. It is Ashby's requisite variety again, and it arrived here from a threshold constant rather than from a ceiling.

Second, it explains why `N -> infinity` is not desirable. As `u_0 -> 0`, any non-zero difference produces unbounded surplus, while `f(0)` remains exactly 0 — the function approaches a step. An observer with infinitely fine discrimination is not maximally perceptive; they are maximally *unstable*, because every difference is total. **The finite threshold is what keeps the lens from having an infinite basement, and finite `N` is what supplies it.**

Third, and this is the payoff: the surplus law was derived in Paper A from four axioms about interaction geometry, with no reference to perception at all. Fechner's law was derived in 1860 from experiments on lifted weights. **They are the same function, and the manual is required to say exactly what that does and does not license.** It does not prove that surplus is perceived, or that perception is surplus. What it shows is that a system built to count difference with diminishing returns and a threshold takes one particular form, and that both the axiomatic route and the psychophysical route arrive at it. **The convergence is real and it is structural. It is not evidence that either is true, and this manual claims nothing further from it.**

36.19 Concavity is the lens, not pessimism

Paper A, Theorem 3.3: `f` is strictly concave in `u`. That is a proved property, unconditional, and it is usually read as bad news — diminishing returns, the fortieth member worth less than the second, effort meeting a wall.

That reading is a coordinate error. §36.18 just showed `f` is a logarithm of `(u + u_0)`, and a logarithm is exactly what a *straight line* looks like when you plot it in the world's coordinates instead of the lens's. Concavity is not the shape of the returns. It is the shape of the lens.

The marginal surplus, `f'(u) = (N−1)/(1+(N−1)u)`, for `N = 12`:

u	f(u)	f'(u)
0.00	0.000000	**11.000000**
0.05	0.438255	7.096774
0.10	0.741937	5.238095
0.20	1.163151	3.437500
0.40	1.686399	2.037037
0.60	2.028148	1.447368
0.80	2.282382	1.122449
1.00	2.484907	**0.916667**

`f'(0) = N − 1 = 11`, which is Paper A's sharp Lipschitz constant (Theorem 4.2) appearing as the initial slope. `f'(1) = 11/12`. **The ratio of first to last marginal value is exactly `N`**, which is a small exact result worth having: `f'(0)/f'(1) = (N−1)/((N−1)/N) = N = 12`, verified.

Now change coordinates. Ask what the world must do to deliver *equal perceived steps* of surplus — 0.5 nats at a time, at `N = 12`:

target F (nats)	required u	(u + u_0)	ratio to previous
0.0000	0.000000	0.090909	—
0.5000	0.058975	0.149884	**1.648721**
1.0000	0.156207	0.247117	**1.648721**
1.5000	0.316517	0.407426	**1.648721**
2.0000	0.580823	0.671732	**1.648721**
2.5000	1.016590 — **impossible**, `u > 1`; ceiling is `ln 12 = 2.484907`		

Every equal step of 0.5 nats multiplies `(u + u_0)` by exactly `e^0.5 = 1.648721`. The "diminishing returns" are a perfectly constant ratio requirement, seen from the wrong side.

Restated as the design principle:

> **Concavity is not a statement that effort stops paying. It is a statement that equal experience costs equal ratios, and equal ratios compound.** The fortieth member is not worth less than the second because members get worse. The fortieth member

is worth less because by then you are asking for a 65 per cent increase in a quantity that is already near its ceiling.

And the ceiling in that last row is worth reading twice: **the fifth equal step is impossible, not expensive.** u would have to exceed 1, and it cannot. `f(1) = ln N = 2.484907` is the hard bound (Paper D, Corollary 2.2). **The lens compresses right up to a wall, and past the wall there is nothing — not diminishing returns but no returns**, which is a different fact and a more important one.

36.20 Doubling the percept: the modality table

`psi = k I^a` doubles when the world quantity is multiplied by `2^(1/a)`. Computed:

Modality	a	world x to **double** the percept	percept of a x2 world	percept of a +10% world
Loudness (1 kHz, sound pressure)	0.67	**2.8138**	x1.5911	+6.59%
Brightness	0.33	**8.1698**	x1.2570	+3.20%
Visual length	1.00	2.0000	x2.0000	+10.00%
Taste (salt)	1.30	1.7044	x2.4623	+13.19%
Heaviness (lifted)	1.45	1.6129	x2.7321	+14.82%
Warmth on arm	1.60	1.5422	x3.0314	+16.47%
Electric shock	3.50	**1.2190**	x11.3137	+39.60%

Two anchors to hold on to. **Doubling apparent brightness takes 8.17 times the luminance. Doubling apparent loudness takes 2.81 times the sound pressure, which is 8.99 dB** — the familiar "about 10 dB per doubling of loudness", recovered from the exponent rather than asserted.

And the last row inverts everything: **doubling the sensation of an electric shock takes only 22 per cent more current.** In the expansive modalities a small world change is a large experience, which is why pain, heat and weight are dangerous to design near and why safety margins in those domains are set in world units rather than perceptual ones.

36.21 Design consequence 1 — budget in ratios

If equal experience is equal ratio, then the unit of a design budget for a perceived quantity is a ratio, not a difference.

Practically: allocate `ln` units. A target of "the room should feel appreciably better" becomes a number — say 0.20 nats of perceived improvement — and that number can be divided across levers and checked by addition, because logs add. This chapter's worked example does exactly this.

The constraint that makes it bite: most world quantities cost money in proportion to the quantity, and deliver experience in proportion to its logarithm. **Cost is linear; experience is logarithmic; therefore the cost of experience is exponential.** That is the fundamental economics of perceived quality and it is why the third refinement always costs more than the first two together.

36.22 Design consequence 2 — never promise a linear improvement in a logarithmic quantity

The failure, quantified. A designer promises "20 per cent brighter" and delivers +20 per cent illuminance:

```
delivered:  1.20 x illuminance
perceived:  1.20^0.33 = 1.06201     ->  +6.201%

the recipient receives 31.0% of what they heard
```

To actually *deliver* a perceived +20 per cent:

Modality	a	world quantity required
Loudness (sound pressure)	0.67	x1.3127 (+31.27%)
Brightness	0.33	**x1.7376 (+73.76%)**
Visual length	1.00	x1.2000 (+20.00%)
Taste (salt)	1.30	x1.1506 (+15.06%)
Heaviness (lifted)	1.45	x1.1340 (+13.40%)
Warmth on arm	1.60	x1.1207 (+12.07%)
Electric shock	3.50	x1.0535 (+5.35%)

A promise of "20 per cent brighter" is a promise to nearly double the lighting load. The promise was made in percent, the delivery is in percent, the two percents are different quantities, and nobody in the conversation noticed. This is the most common and most expensive error the lens produces, and the correction is a rule rather than a calculation:

> **RULE.** State the promise in the recipient's units and convert to yours before agreeing to it. **A commitment expressed as a percentage of a world quantity is not a commitment about experience, and must not be sold as one.**

36.23 Design consequence 3 — geometric ladders

Anything intended to deliver equal steps of experience must be a geometric progression. A five-rung ladder from 40 000 to 100 000 with equal perceived steps:

```
common ratio = (100000/40000)^(1/4) = 1.257433      (+25.74% per rung)
```

rung	value	ln
0	40 000.00	10.596635
1	50 297.34	10.825707
2	63 245.55	11.054780
3	79 527.07	11.283853
4	100 000.00	11.512925

Each `ln` step is 0.229073, identical. Compare the linear ladder over the same range, in equal absolute steps of 15 000:

rung	value	ratio to previous	ln step
0	40 000.00	—	—
1	55 000.00	1.375000 (+37.50%)	0.318454
2	70 000.00	1.272727 (+27.27%)	0.241162
3	85 000.00	1.214286 (+21.43%)	0.194156
4	100 000.00	1.176471 (+17.65%)	0.162519

The linear ladder's perceived steps shrink monotonically: the last step is barely half the first. An equal-increment ladder is experienced as a decelerating one, and the person on the top rung feels the least rewarded.

This manual is required to note the limit of the example: **whether money is perceived logarithmically is a contested empirical question and this chapter does not settle it.** The arithmetic above is a demonstration of what a geometric ladder does to *any* logarithmically perceived quantity; whether compensation is one is a matter for measurement, and the same construction applies without argument to quantities that are known to be logarithmic — light, sound, and the surplus law of §36.18.

36.24 Design consequence 4 — report geometric means and log axes

Three reporting rules, each a direct corollary of a table row:

1. **Central tendency of a perceived quantity is the geometric mean** (§36.10). Report it. Report the arithmetic mean only when the quantity is genuinely additive in the world and no one is experiencing it.
2. **Plot on a log axis.** A log axis is the lens. A quantity plotted linearly and perceived logarithmically produces a chart whose visual impression contradicts the experience it describes.
3. **Report ratios, not differences.** "From 200 to 300 lux" is +50 per cent and 0.405465 nats of world; "an increase of 100 lux" is a number with no perceptual content until you know where it started.

36.25 Design consequence 5 — the reference is cheaper than the quantity

From §36.13: the percept is `k ln(I/I_0)` and `I_0` is set by adaptation and context. **Moving the reference by a factor of two has exactly the same perceptual effect as moving the quantity by a factor of two, and it is usually far cheaper.**

This is not manipulation and it is not free of ethics; §36.27 addresses the boundary. It is a real engineering lever, it is what a dim vestibule before a bright room is, and it is what a moment of silence before a soft sound is. **The designer who only ever moves `I` is working with one of the two available terms.**

36.26 Where the lens does not apply

Required by the honesty clauses, and by CANON §5.

- **Expansive modalities invert the advice.** At `a > 1` the world compresses rather than the percept, and a small world increment is a large experience. Every rule above about needing large ratios reverses. **Look the exponent up. Do not assume compression.**

- **Small counts are not logarithmic.** Numerosity for very small sets is direct and near-linear. The lens is a law about magnitude, not about counting to four.
- **Visual length is essentially linear** (a ~ 1.0). For spatial extent, differences behave like differences, and this chapter's corrections are unnecessary and would be wrong.
- **The lens describes magnitude judgment, not value, preference or decision.** Extending it to those is a separate empirical claim requiring separate evidence, and this manual does not make it.
- **Outside the working range, Weber's fraction is not constant.** Near threshold and near saturation the law degrades, which is documented and is a known limit rather than a hidden one.
- **The constants are individual.** The form recovers across people; c, k and a vary. Where precision matters, measure the person.

36.27 The ethical boundary

A lens argument can be used to give someone more experience for the same resources, or to give them the same experience for fewer resources while telling them it is more. The arithmetic is identical and the two uses are not.

CANON §5.4 governs and the operative test is disclosure and decline. **A reference-shifting technique that the recipient would object to if told about it is one this manual does not endorse, and the fact that it is efficient does not change that.** The vestibule that makes the reading room feel bright is disclosed by walking through it. A brightness claim computed in world units and sold in perceptual ones is not.

The Specification

S36.1 — Declare the register

Any quantity used in design **shall** be labelled as perceived logarithmically, perceived linearly, or unknown. **Where it is unknown, no perceptual claim shall be made about it.**

S36.2 — Convert before promising

Any commitment about experience **shall** be converted to world units using a stated exponent before it is agreed, and the exponent and its source **shall** be recorded.

S36.3 — Ratios in budgets, never differences

Design budgets for perceived quantities **shall** be stated as ratios or in log units. **A budget stated as an absolute difference in a logarithmically perceived quantity shall be rejected as unspecified.**

S36.4 — Geometric mean for central tendency

Reports of a perceived quantity **shall** give the geometric mean. **The arithmetic mean of ratios shall not be reported** (see also S35.2 — it is the same rule).

S36.5 — Log axes

Charts of perceived quantities **shall** use logarithmic axes, labelled as such.

S36.6 — Equal-step ladders shall be geometric

Any schedule intended to deliver equal perceptual increments **shall** be a geometric progression, and the common ratio **shall** be stated.

S36.7 — Exponents cited with their method

Any Stevens exponent used **shall** be reported with its modality, its stimulus range, and the fact that it is a constant of a procedure. **An exponent quoted bare shall be treated as unsourced.**

S36.8 — Do not divide Fechnerian percepts

Where the percept scale is Fechnerian (interval), **only differences shall be taken.** Ratios of percepts require a magnitude-estimation scale and **shall** state that they have one.

S36.9 — Do not sum across modalities without a declared weighting

Perceived changes in different modalities are in different units. **They shall not be added without an explicit, stated exchange rate**, and the exchange rate shall be badged as a design decision rather than a measurement.

S36.10 — Threshold declared

Any use of a logarithmic model **shall** state the reference `I_0` or threshold `u_0` in use. **A logarithm with an undeclared reference is not a measurement.**

S36.11 — Reference shifts disclosed

Where a design achieves its effect by moving the reference rather than the quantity, that **shall** be disclosable to the occupant without embarrassment. **If it could not be, it shall not be built.**

S36.12 — Concavity reported as a ratio requirement

Where a proposal is rejected on grounds of diminishing returns, the rejection **shall** state the ratio increment that would have been required. **"Diminishing returns" without a number is not a finding.**

The Failure Modes

FM-36.1 — The Linear Promise

Description. A commitment about experience is made in world percentages. "Twenty per cent brighter", "twice as fast", "half the noise."

Damages. The relationship, and the budget. In brightness, +20 per cent delivered is +6.2 per cent received — 31 per cent of the promise.

Diagnostic signature. A percentage in a specification with no exponent anywhere in the document.

Fix. S36.2. Convert before agreeing. **The conversion takes one line and it is the highest-value single line in this chapter.**

FM-36.2 — The Arithmetic Mean of Ratios

Description. Percentage changes or ratios are averaged arithmetically.

Damages. The number, always in the same direction (overstatement, by Jensen).

Diagnostic signature. +50 per cent and -50 per cent averaging to zero when the true per-step factor is 0.866025, that is -13.40 per cent.

Fix. S36.4. Geometric mean, or mean of logs.

FM-36.3 — The Linear Ladder

Description. Equal absolute increments — of pay, of capacity, of allowance — intended to feel equal.

Damages. The perceived step decays monotonically. In the §36.23 example, the last step is 0.1625 nats against the first's 0.3185, barely half.

Diagnostic signature. Compute the ln steps of any ladder. Unequal ln steps mean unequal experience.

Fix. S36.6. Geometric progression with a stated common ratio.

FM-36.4 — Diminishing Returns Read as Failure

Description. A programme delivering constant ratio improvement is reported as decelerating and cancelled.

Damages. The programme, and the credibility of the next one.

Diagnostic signature. Plot the log. **If the log is a straight line, the programme is delivering constantly and the chart is lying.**

Fix. S36.5 and S36.12. Plot on a log axis; state the ratio requirement rather than the word "diminishing".

FM-36.5 — The Undeclared Reference

Description. A logarithm is taken with no stated reference or threshold. The resulting numbers are offsets from an unknown origin.

Damages. Comparability. Two teams' numbers cannot be combined and neither knows it.

Diagnostic signature. Ask what value gives zero. If nobody can say, S36.10 has failed.

Fix. Declare `I_0`. In the surplus law, declare `N`, since `u_0 = 1/(N−1)`.

FM-36.6 — Nats Added Across Modalities

Description. Perceived improvements in lighting, acoustics and thermal comfort are summed into a single figure.

Damages. The figure. Nats of brightness are not nats of loudness.

Diagnostic signature. A single "perceived improvement" number covering more than one sense, with no exchange rate stated.

Fix. S36.9. Report per modality. If a single figure is genuinely needed, state the weighting as a design decision and badge it.

FM-36.7 — The Dimensioned Logarithm

Description. `ln(I)` computed where `I` carries units — `ln(500 lux)`.

Damages. Dimensional coherence. FVP-1 fails.

Diagnostic signature. A log whose argument has a unit attached.

Fix. Normalize first: `ln(I/I_0)`. Always.

FM-36.8 — The Expansive Modality Treated as Compressive

Description. Chapter rules about needing large ratios are applied to heat, weight, or pain.

Damages. Safety, potentially. At `a = 3.5`, a 22 per cent increase doubles the sensation, and a design margin computed on compressive assumptions is far too small.

Diagnostic signature. No exponent looked up.

Fix. S36.7. Look it up. **The sign of the design advice depends on whether a is above or below 1, and there is no default.**

FM-36.9 — Percepts Divided

Description. Two Fechnerian percept values are divided and reported as "twice as much".

Damages. The claim. An interval scale has no true zero and does not support division.

Diagnostic signature. A ratio of percepts with no magnitude-estimation procedure behind it.

Fix. S36.8. Either take differences, or build a ratio scale by magnitude estimation and say that you did.

FM-36.10 — The Reference Shift Sold as a Quantity Increase

Description. Experience is improved by moving `I_0` and reported as an improvement in `I`.

Damages. Honesty, and the next measurement, which will not reproduce.

Diagnostic signature. Reported world quantity unchanged; occupant reports improved; report claims the quantity improved.

Fix. S36.11. Report the mechanism. **Reference shifting is a legitimate technique and a dishonest claim, and the difference is entirely in the reporting.**

Worked Example

The reading room, budgeted in ratios

A public library reading room, treated in Volume One as a circuit (Chapter 12), returns here as a perceptual problem. Three complaints and one constraint.

```
  Room:        12.0 m x 8.0 m x 3.5 m,  V = 336.0 m^3
Complaint 1: "the edges are gloomy"     perimeter 200 lux, core tables 500 lux
Complaint 2: "you can hear everything"   8 simultaneous quiet talkers; RT60 measured 1.2
Complaint 3: "it feels harsh"          (deferred — not a magnitude problem)
  Constraint:  lighting energy must not rise more than 10 per cent
```

Step 1 — Convert every complaint into a ratio. No absolute numbers are actionable until they are ratios against a reference.

Gloom. Perimeter against core:

```
  world ratio   200/500 = 0.400000        ln = -0.916291 nats
perceived    (0.400)^0.33 = 0.7391    -> the perimeter looks 26.09% dimmer than the core
  in percept units: a . ln(ratio) = 0.33 x (-0.916291) = -0.302376 nats
```

A 60 per cent reduction in illuminance is experienced as a 26 per cent reduction in brightness. The lens has already absorbed most of the physical difference, which is why the complaint is "gloomy" rather than "dark" — and it is also why the fix is going to be more expensive than the complaint sounds.

Noise. Eight incoherent talkers against a target of four. Incoherent sources add in power, so halving the count divides pressure by `sqrt(2)`:

```
pressure ratio  1/sqrt(2) = 0.707107     = -3.010 dB
perceived       (0.707107)^0.67 = 0.7928 -> 20.72% quieter
in percept units: 0.67 x ln(0.707107) = -0.232204 nats
```

Halving the number of people talking makes the room about a fifth quieter, not half. This single number resets the expectations of every occupancy-reduction proposal the library has considered.

Reverberation. Sabine: `RT60 = 0.161 V / A`.

```
measured RT60 = 1.20 s   =>   A = 0.161 x 336.0 / 1.20 = 45.08 m^2 sabins
target  RT60 = 0.60 s    =>   A = 0.161 x 336.0 / 0.60 = 90.16 m^2 sabins
 required: A x 2.000  -- exactly double the absorption
```

Halving reverberation time requires doubling absorption, exactly, because Sabine's law is itself a reciprocal — a ratio law. The perceptual conversion stops here and this manual stops with it:

NOT CONVERTED. There is no exponent in this manual's verified table for perceived reverberance. The world ratio is `ln(0.5) = −0.693147` nats and **it shall not be multiplied by an exponent that has not been measured.** The acoustic change is reported in world units and left there. S36.7 forbids the alternative, which would have been to invent a plausible number.

Step 2 — Price the constraint before designing to it. The energy cap is +10 per cent. What does the lens permit inside it?

illuminance change	energy change	perceived brightness change
500 -> 450 lux	-10.0%	**-3.42%**
500 -> 400 lux	-20.0%	**-7.10%**
500 -> 350 lux	-30.0%	-11.10%
500 -> 250 lux	-50.0%	-20.45%
200 -> 300 lux	+50.0% (perimeter circuit)	**+14.32%**
200 -> 400 lux	+100.0% (perimeter circuit)	+25.70%
500 -> 4 085 lux	+717%	+100% (a perceived doubling)

The bottom row is the anchor. **Doubling the apparent brightness of the core would take 4 085 lux and eight times the energy. It is not on the table and never was.** Knowing that before the meeting is what a table like this is for.

The fourth row is the one that solves the problem. **A 20 per cent energy cut on the core costs 7.1 per cent of perceived brightness**, and the fifth row spends that saving where it buys three times as much: **a 50 per cent increase on the perimeter buys 14.3 per cent of perceived brightness there, because the perimeter starts lower and the lens is steeper in ratio terms near the bottom of the range.**

Step 3 — The design, and its energy arithmetic. With the core carrying 60 per cent of the luminaires and the perimeter 40 per cent:

```
core:      500 -> 400 lux    (x0.80)
perimeter: 200 -> 300 lux    (x1.50)

energy = 0.60 x 0.80 + 0.40 x 1.50 = 0.4800 + 0.6000 = 1.0800    -> +8.00%
```

Inside the 10 per cent cap.

Perceptual result:

location	before	after	perceived change
Core tables	500 lux	400 lux	-7.10%
Perimeter	200 lux	300 lux	**+14.32%**
Core-to-perimeter gap	-26.09%	**-9.06%**	gap cut by 65%

The complaint was the gap, and the gap closed from 26 per cent to 9 per cent — for 8 per cent more energy — by taking brightness away from the place that had the most of it. In world units that move looks perverse: you reduced the light at the reading tables. In lens units

it is straightforwardly correct, because 100 lux removed from 500 costs 7 per cent of experience and 100 lux added to 200 buys 14 per cent. **The lens is why the two 100-lux moves are not the same move.**

Step 4 — The perceived-improvement budget, per modality. In nats, with k = 1 and the modality's exponent as the gain, and **kept strictly separate by modality as S36.9 requires**:

```
BRIGHTNESS (a = 0.33)
   core      500 -> 400 lux   : 0.33 x ln(0.800000) = -0.073637 nats   (-7.10%)
   perimeter  200 -> 300 lux   : 0.33 x ln(1.500000) = +0.133803 nats  (+14.32%)
    gap closure (the design objective): from -0.302376 to -0.094839 nats

LOUDNESS (a = 0.67)
   8 -> 4 talkers            : 0.67 x ln(0.707107) = -0.232204 nats  (-20.72%)
   -- a REDUCTION in loudness, which is an improvement; sign is in the modality, not the ledg

REVERBERATION
    RT60 1.20 -> 0.60 s         : world ln-ratio -0.693147 nats
                                  ** NOT CONVERTED - no verified exponent **
```

These three blocks are not summed and this manual will not sum them. A single "perceived improvement of 0.44 nats" would require an exchange rate between brightness nats and loudness nats, no such rate has been measured here, and inventing one is FM-36.6. The report gives three figures and says so.

Step 5 — Where the acoustic budget actually goes. The library's instinct was occupancy limits. Priced:

intervention	world change	perceived loudness
8 -> 4 talkers	-3.010 dB	-20.72%
Pressure halved (-6.02 dB)	x0.5 pressure	-37.15%
-9.0 dB	x0.3548 pressure	**-50.05%**
-10.0 dB	x0.3162 pressure	-53.76%
-12.0 dB	x0.2512 pressure	-60.37%

Halving the perceived loudness of the room takes -8.99 dB, which is a factor of 0.3554 in pressure — the equivalent of going from eight talkers to one. Occupancy control cannot deliver it; the room would have to be nearly empty. The absorption route can: doubling A to 90.16 m² of sabins halves RT60, which reduces the reverberant contribution to the level at every listening position, and it does so without removing a single reader.

That is the finding the arithmetic produced and the meeting would not have. **The library's plan was to reduce the number of people in a reading room; the lens says the achievable version of that plan delivers a fifth of what the complaint is asking for, and the ceiling and wall treatment delivers the rest without touching occupancy at all.**

Step 6 — What to promise. S36.2 requires the commitment be stated in the recipient's units and converted before agreement. The library's draft notice said *"we are increasing lighting at the perimeter by 50 per cent."*

Converted: readers will experience **about 14 per cent** more brightness at the perimeter, and about 7 per cent less at the tables, and the difference between the two zones — which is what they complained about — will be **about a third of what it was**.

The revised notice says that. It is a smaller number and a truer one, and it is the number the building will actually deliver. **A promise of 50 per cent, kept exactly in lux, would have been received as a failure, because the promise was made in one register and the reading is taken in another.** That is the Lens Theorem's entire practical content, and it has now cost the library nothing except the willingness to compute a logarithm before writing a notice.

Step 7 — What would falsify the design's predictions. The library preregisters, before the work: post-occupancy magnitude estimation (Chapter 38's method) at the perimeter and at the core, by readers who did not participate in the design, against the prediction of **+14.32 per cent and -7.10 per cent respectively**, with the gap closing to -9.06 per cent. **If the measured perimeter improvement is outside a factor corresponding to ±0.05 nats of the prediction, the brightness exponent of 0.33 does not describe this population in this room, and the exponent — not the design method — is what gets revised.**

That distinction is worth ending on. **The Lens Theorem is not at risk from this measurement, because it is a theorem. The exponent is at risk, and it should be, because it is a measurement.** Knowing which of your commitments are proofs and which are measurements is the difference between a manual and a manifesto, and it is the discipline this volume's remaining chapters are built to enforce.

Chapter 37 turns to the methodology this chapter has been quietly demonstrating: how a subjective structure can be postulated, have its consequences derived, and be held to them — which is exactly how psychophysics succeeded, and exactly what Chapters 34 and 35 have been asking of the Moment Equation.

Chapter 37 . Subjective Reality as Postulate

Volume One measured rooms: turn-taking, reverberation time, communication paths, turnover, decision latency. Every one of those quantities has the convenient property that a second observer can walk in with a stopwatch and a sound level meter and get the same number. Volume Two has now spent five chapters proposing that the same instrument works on the inside of a single head, and a reader who has been paying attention has been holding one objection in reserve since Chapter 32.

The objection is this. **You cannot walk into someone's experience with a stopwatch.** There is no second observer. There is no instrument that touches the quantity directly. Every number in Part VII is therefore either a report about an inner state, which is not the state, or a behavioural proxy for an inner state, which is also not the state. If the quantity cannot be observed, the formulas that operate on it are decoration.

This chapter answers that objection by naming the method the sciences of the inside actually use, stating the conditions under which it is legitimate, and stating the conditions under which it degenerates into the thing the objection accuses it of being.

The method is: you cannot measure the inner state, so you postulate a structure for it, you derive consequences that are observable, and you test the consequences. The postulate is never tested directly. It is tested through what it forbids. If the forbidden things happen, the postulate is dead. If they do not, the postulate has survived a round and is still on probation, which is the only condition any postulate is ever in.

This is the methodological keystone of Volume Two. Chapter 38 will show that this exact method produced a century of replicable ratio-scale measurement of subjective magnitude. Chapters 39 through 51 will use it on framings, moments, gradients and echoes. Chapters 52 through 56 will harden it into a protocol. **Everything downstream of here is licensed by this chapter or by nothing.**

The Law

37.1 What is proven here, and what is not

Nothing in this chapter is a theorem. There is no Paper A result about postulation, because postulation is not a mathematical object — it is a discipline of practice. What this chapter offers instead is: a precise statement of the **Postulate Method**, badged as a house formalisation of standard hypothetico-deductive practice; six **admissibility conditions** separating a legitimate postulate from an unfalsifiable one; one genuine derivation — **Fechner's integration of Weber's law** — as the exemplar, because it is the historical case in which the method worked on exactly this problem; two quantitative instruments, an **empirical content measure** and a **held-out skill score**, each with its full FVP table; and the **badge convention** itself.

Everything asserted about the world in this chapter is asserted about the *practice of inquiry*, not about minds. The claims about minds start in Chapter 38 and they arrive with data.

37.2 The Postulate Method, stated

> **THE POSTULATE METHOD [HOUSE FORMALISATION].** Let X be a quantity you cannot observe. Then:
>
> **Step 1 — POSTULATE.** Assert a structure for X: what kind of thing it is, what operations it supports, how it relates to quantities you *can* observe. State the assertion in the open, with its badge on, in the same sentence in which you first use it.
>
> **Step 2 — DERIVE.** From the postulate plus already-accepted results, derive statements about observables. The derivation must be valid; if it is not, nothing that follows means anything. A derived statement counts only if the negation of the postulate would not also produce it.
>
> **Step 3 — FORBID.** Extract from the derived statements the set of observations the postulate rules out. Write them down before you look. This is the postulate's content and it is the only part of it that can be tested.
>
> **Step 4 — TEST.** Observe. If a forbidden observation occurs, the postulate is dead as stated. If none occurs, the postulate has survived, which is not the same as being true.
>
> **Step 5 — DECLARE.** Publish the surviving postulate *with its badge still on*. Survival does not promote a postulate to a theorem. Nothing promotes a postulate to a theorem.

Read Step 2's second sentence again; it is where most work in the human sciences fails. **A derivation that both the postulate and its negation produce is not evidence for the postulate.** If "I am a morning person" predicts that some mornings go well, and its negation predicts the same, then mornings going well is not data about the postulate. It is data about mornings.

37.3 The badge taxonomy — the house convention in full

Volume One introduced badging informally: *the mysticism is permitted, but it is permitted in the open, with its badge on* (CANON §2). Volume Two needs the badges to be a closed, small, non-overlapping set, because from here on almost every sentence in the book is one of these things and the reader is entitled to know which.

Badge	Meaning	Test that applies	Can it be wrong?
[PROVEN]	A theorem. Follows from stated axioms by valid deduction.	Check the proof.	Only if the proof is invalid.
[DESIGN AXIOM]	Chosen, not derived. Fixes a modelling convention.	None. It is a decision.	It can be *unwise*, not false.
[POSTULATE]	An assertion about the world that generates forbidden observations.	Steps 3–4 above.	**Yes, and how is stated.**

Badge	Meaning	Test that applies	Can it be wrong?
[HEURISTIC]	A rule of thumb that orders or triages. Not an equation.	Does it improve triage?	It can be useless.
[INTERPRETIVE CORRESPONDENCE]	Two formalisms have the same shape. No theorem connects them.	None available.	It is not the kind of thing that is true or false.
[METAPHOR]	A figure of speech carried deliberately.	None.	No. It carries no claim.

Four rules govern the use of the table.

1. **The badge appears in the same sentence as the first statement of the thing.** Not in a footnote, not in an appendix, not "as discussed earlier".
2. **A badge never upgrades through use.** A postulate repeated for two hundred pages is still a postulate on page two hundred. This is the commonest corruption in the literature of the inside: an assumption made once, referenced, then referenced by the reference, until the fourth-generation citation reads as a finding.
3. **A [POSTULATE] that cannot state what would kill it is not a [POSTULATE].** It is at best a [HEURISTIC] and must be demoted in print, visibly.
4. **[INTERPRETIVE CORRESPONDENCE] and [METAPHOR] carry no evidential weight whatever.** They may be beautiful and productive. They may not be cited as support for anything.

The badge is not a hedge. It is a specification of what kind of test the reader is entitled to demand. [PROVEN] licenses a demand for a proof; [POSTULATE] for the list of forbidden observations; [DESIGN AXIOM] for a justification of the choice, and for a counter-proposal. The badge tells the reader which demand is in-bounds. That is a courtesy, not a defence.

37.4 Six admissibility conditions for a postulate

A postulate is admissible in this manual if and only if it satisfies all six.

```
P1 OPENNESS     It is stated explicitly, in the open, badged, before it is used.

P2 ASYMMETRY    It generates at least one observable statement that its own negation
                does not generate. (Step 2's second sentence, made a gate.)

P3 REACHABILITY The forbidden observations are reachable with instruments the reader
                   actually has, at a cost the reader can actually pay.

P4 LETHALITY    A specific, named, pre-recorded result kills it. Written before looking.

P5 MINIMALITY   It does not smuggle. Every clause does work; a clause that forbids
                   nothing must be deleted, not carried.

P6 STABILITY    Its parameters recover across scopes and observers, or the postulate
```

```
                states plainly that it is indexed to one observer and claims nothing wider.
```

P1 through P4 are the standard falsificationist demands, sharpened. P5 and P6 are the two that most work in the human sciences fails.

P5 (minimality) is a defence against the protective belt. Lakatos (1970) described how a research programme protects its core by accreting auxiliary hypotheses, each added to absorb a disconfirmation. Each addition is locally reasonable; the aggregate is a postulate that forbids nothing, because whatever happens one of the clauses accommodates it. A postulate should get *smaller* as it survives, not larger.

P6 (stability) is FVP-7 in advance. A postulate about *your* inner structure that holds only for you is admissible — single-subject science is real science — **but it must say so, and it must not be quoted as though it were a claim about people.** Most confident statements about "how the mind works" are single-subject claims about the speaker, generalised without the badge. Chapter 54 makes this gate formal via specific objectivity; here it is a discipline of speech.

37.5 The exemplar: how Fechner did it

The reason to believe the Postulate Method works on inner states is not an argument. It is a case. In 1860 Gustav Fechner published *Elemente der Psychophysik* and did precisely Steps 1 through 5 on the least tractable quantity available — the magnitude of a sensation — and the result is still in use.

He began with something measured. Ernst Weber (1834) had reported that the smallest detectable change in a stimulus is a constant *proportion* of it, not a constant amount:

```
    WEBER'S LAW  [EMPIRICAL REGULARITY, measured]

       dI / I = c                    c is the Weber fraction for the modality

    A weight of 100 g needs about the same fractional increment to feel heavier as a
    weight of 1000 g does. The increment in grams is ten times larger; the ratio is
    the same. That is the finding, and it is a finding about behaviour: what people
     can and cannot discriminate. Nothing about inner magnitude is asserted yet.
```

Weber's law is on the observable side of the wall. It concerns discrimination performance, which you can test with a balance and a protocol. Fechner wanted the quantity on the other side of the wall — the felt intensity itself — and he could not reach it. So he postulated a bridge:

> **FECHNER'S POSTULATE [POSTULATE, 1860].** All just-noticeable differences are subjectively equal. The step in sensation produced by one JND is the same size wherever on the scale it is taken.

That is the whole move, and it is worth pausing on how audacious and how modest it is at once. It is audacious because it asserts an equality between inner quantities that no one can inspect. It is modest because it asserts exactly one thing and immediately hands the consequences to arithmetic. Combine it with Weber and integrate:

```
    THE FECHNER DERIVATION

      postulate:  dS = k' * (one JND)        all JNDs equal in sensation units
```

```
Weber:       one JND = c * I              in stimulus units
therefore:   dS = k * dI / I              k absorbs k' and c

integrate:   S = k * ln(I) + const
set S = 0 at absolute threshold I_0:

    S = k * ln( I / I_0 )                 FECHNER'S LAW
```

Verify the ladder arithmetic. With a Weber fraction c = 0.05, the stimulus after n just-noticeable steps is `I_n = I_0 (1+c)^n`, and inverting recovers the count exactly:

```
TABLE 37.1 — THE JND LADDER, c = 0.05 (computed)

     n        I_n            ln(I_n/I_0) / ln(1+c)
   ---    ----------        ---------------------
     0          1.0000                    0.0000
    10          1.6289                   10.0000
    50         11.4674                   50.0000
   100        131.5013                  100.0000
   200      17292.5808                  200.0000
   283     992136.9785                  283.0000

  Sensation measured in JND units is exactly the log of the stimulus ratio divided
   by ln(1+c).  The constant is k = 1 / ln(1.05) = 20.495934 JNDs per nat.
```

Now the crucial step, the one that makes this an exemplar rather than an anecdote. **Fechner's postulate is not tested by looking at it. It is tested by its consequences, which live entirely on the observable side of the wall.** The law forbids things: it forbids equal stimulus *differences* producing equal report differences across the range; it forbids a modality with a constant Weber fraction from showing a linear relation between stimulus and any well-constructed sensation scale; it forbids the JND count between two stimuli from depending on the path taken between them; and it predicts, quantitatively, the number of discriminable steps across a modality's dynamic range.

Every one of those is testable with nineteenth-century equipment. Some of them turned out to be wrong, which is the point: Chapter 38 documents how Stevens (1957) killed Fechner's law as a general account and replaced it with a power law, using a method Fechner had not thought of. **The postulate died the death a legitimate postulate is supposed to die: by its own consequences, at the hands of a better instrument.** A working postulate is not one that survives forever. It is one that fails informatively.

37.6 The FVP table for the exemplar

Every formula in Volume Two carries its seven gates (CANON_INTERNAL §6). Fechner's law goes first because it is the template.

```
FVP — FECHNER'S LAW    S = k ln(I / I_0)
```

Gate	Verdict for Fechner's law
FVP-1 DIMENSIONS	`I/I_0` is a dimensionless ratio, so the logarithm is well formed. `S` inherits the units of `k`, which are *sensation units per nat* — in the derivation above, JNDs per nat, numerically `1/ln(1+c)`. Both sides agree. **The dimensional health of the law depends entirely on `I_0` being a real quantity in the same units as `I`**, i.e. an absolute threshold, not a convenience constant.
FVP-2 SCALE TYPE	`I` is ratio-scaled (physical intensity, true zero). `S` as constructed is **interval-scaled at best**: the JND-counting procedure fixes a unit and an origin by convention, and Fechner's postulate is precisely the assertion that the unit is constant. It does not establish that `S` has a true zero, so **ratios of Fechner sensations are not licensed**. This is a real limitation and it is the gap Stevens' magnitude estimation later closed by a different route (Chapter 38, Chapter 53).
FVP-3 LIMITING CASES	`I -> I_0`: `S -> 0`, sensation vanishes at threshold — correct and required. `I -> 0`: `S -> -infinity`, which is *not* a real phenomenon and correctly signals that the law is undefined below absolute threshold; the domain must be restricted to `I >= I_0`. `I -> infinity`: `S -> infinity` unboundedly, which is wrong — real modalities saturate, and at high intensity the tissue is damaged rather than the sensation extended. **The law is a mid-range law and its own limits announce this.**
FVP-4 FIXED POINTS	`S = 0` at `I = I_0` is the only fixed structure. The scale-invariance property is the substantive invariant: `S(aI) - S(I) = k ln a` for any a, **independent of `I`**. Equal ratios produce equal sensation differences everywhere. That invariance is the law's whole content and it is what Chapter 36's Lens Theorem is built on.

Gate	Verdict for Fechner's law
FVP-5 IDENTIFIABILITY	k and `I_0` are jointly recoverable from discrimination data: `I_0` from threshold measurement, k from the Weber fraction as `1/ln(1+c)` in JND units. **They are not separately identifiable from magnitude ratings alone**, because a rating scale with arbitrary origin and unit absorbs both. This is a genuine identifiability failure and it is why the law must be fitted to discrimination data, not to ratings.
FVP-6 FALSIFICATION	Killed by: a modality with a stable Weber fraction whose independently constructed sensation scale is not logarithmic in intensity. Stevens' magnitude estimation supplied exactly this for several modalities, and additionally showed exponents above 1 (heaviness, warmth, shock) which no logarithmic law can produce. **Fechner's law is falsified as a general law and survives as an approximation over restricted ranges** — see Chapter 38's computed comparison.
FVP-7 REPRODUCIBILITY	Weber fractions and absolute thresholds recover across observers and laboratories to within modality-typical scatter; this gate is passed for the *inputs*. The postulate that JNDs are subjectively equal is not itself reproducible-or-not, because it is not measurable — **which is exactly why it is a postulate and exactly why it had to be tested through its consequences.**

Look at the FVP-7 row. It states plainly that the load-bearing postulate is untestable in isolation, and it does so without embarrassment, because the method never required it to be testable in isolation. That row is the model for every FVP-7 in this volume.

37.7 The honesty clauses for this chapter

1. The Postulate Method is a formalisation of ordinary hypothetico-deductive practice. **It is not original to this manual and no priority is claimed.**
2. The Duhem–Quine problem applies in full: when a forbidden observation occurs, logic alone does not say whether the postulate failed or an auxiliary assumption did. **The method does not dissolve this; it manages it by requiring the auxiliaries to be listed in advance**, so the choice of what to blame is made in the open.
3. Survival is not confirmation. The manual never writes "confirmed". The permitted words are *survived*, *not yet killed*, *consistent with*.

4. Nothing here licenses a claim about consciousness, physics, or thermodynamics (CANON §5.1).

The Engineering

37.8 The four disguises of the unfalsifiable postulate

An unfalsifiable postulate is rarely announced as one. It is usually a legitimate postulate that has been damaged in one of four specific ways, each with a clean diagnostic.

Disguise 1 — THE UNIVERSAL PREDICTOR. The postulate is compatible with every outcome. Its signature is that you cannot complete the sentence "if I observe _____, I will abandon this." Diagnostic: hand the postulate to someone hostile and ask them to design an experiment to kill it. If they cannot, and they are competent and motivated, the postulate has no content.

Disguise 2 — THE ELASTIC TERM. The postulate is falsifiable in principle but one of its terms is defined loosely enough to stretch over any disconfirmation. "I work best under pressure" fails when a pressured week goes badly, unless "pressure" quietly becomes "the right kind of pressure", which is the pressure of the weeks that went well. Diagnostic: **the operational definition must be written down before the observation and must be executable by someone who does not know the hypothesis.** If your definition requires your judgement to apply, it is elastic.

Disguise 3 — THE POST HOC FIT. The structure is proposed after the data are in and is therefore guaranteed to fit them. This is not fraud; it is the normal and productive activity of generating hypotheses, and it becomes a fault only when the fit is reported as a test. Diagnostic: a date. If there is none, treat it as post hoc, including your own.

Disguise 4 — THE RETREAT. Each time a forbidden observation occurs, a clause is added. The postulate never dies; it accumulates, and its content shrinks toward zero while its word count grows. Diagnostic: §37.9's amendment ledger, which makes the shrinkage visible.

37.9 Legitimate amendment versus retreat

Postulates are supposed to be amended; amendment is how a research programme learns. The distinction between amendment and retreat can be made precise.

> **THE AMENDMENT RULE [HOUSE RULE].** An amendment to a postulate is legitimate if and only if the amended postulate forbids at least one observation that the original did not, and the new forbiddance is recorded, dated, and tested in a *subsequent* round on *fresh* data. An amendment that only removes a forbiddance is a retreat and must be recorded as one.

The instrument that enforces the rule is a ledger. Every postulate you carry gets one.

```
AMENDMENT LEDGER — one row per revision

date | version | clause added or removed | forbiddance ADDED | forbiddance REMOVED | net
```

A programme whose "net" column is negative for three consecutive revisions is degenerating, and the honest move is to abandon the postulate rather than revise it a fourth time. This

is Lakatos' progressive/degenerating distinction rendered as bookkeeping. It must be a written artifact and not an intention: an artifact has low delta (Chapter 3), an intention has delta near 1.

37.10 Empirical content: measuring what a postulate forbids

Popper's central insight was that the content of a theory grows with the class of what it forbids — a theory compatible with everything says nothing. That insight can be given a number, and the number lands in the same logarithmic register as everything else in this volume, which is convenient but is not evidence for anything.

> **THE CONTENT MEASURE [HEURISTIC].** Let a postulate admit a fraction q of the outcomes you considered possible before adopting it, 0 < q <= 1. Define its empirical content as
>
> ```
> H_forbid = −ln q          nats
> ```
>
> q = 1 (forbids nothing) gives zero content. Halving the admitted outcomes adds ln 2 nats.

TABLE 37.2 — EMPIRICAL CONTENT OF A POSTULATE (computed)

q (fraction of outcomes still admitted)	nats	bits
1.000	0.0000	0.0000
0.900	0.1054	0.1520
0.500	0.6931	1.0000
0.250	1.3863	2.0000
0.100	2.3026	3.3219
0.050	2.9957	4.3219
0.010	4.6052	6.6439
0.001	6.9078	9.9658

The table is used comparatively, never absolutely. Two postulates about the same question, over the same enumerated outcome set, can be ranked by content, and **the higher-content postulate is the better one to test first.** A postulate at q = 0.9 costs a full experiment and, if it survives, has bought 0.105 nats. A postulate at q = 0.05 costs the same experiment and buys 2.996 nats — nearly thirty times the return for the same expenditure. That ratio is the whole practical argument for preferring bold postulates to safe ones, and it is arithmetic rather than temperament.

FVP — THE CONTENT MEASURE H_forbid = −ln q

Gate	Verdict
FVP-1 DIMENSIONS	q is a dimensionless fraction in (0,1]. H_forbid is in nats. Dimensionally sound. **Its meaningfulness, however, depends on the outcome set being enumerated and fixed in advance** — q is a fraction *of something*, and if that something is chosen after the fact the number is meaningless.

Gate	Verdict
FVP-2 SCALE TYPE	q requires a ratio scale on the outcome space, i.e. a defensible measure over outcomes. In practice you almost never have one. **Therefore this is a [HEURISTIC] and not an equation, and it is demoted in print accordingly**: use it to rank two postulates over a shared, explicitly listed outcome set, never to report an absolute content.
FVP-3 LIMITING CASES	`q -> 1`: content `-> 0`. Correct: a postulate that forbids nothing says nothing. `q -> 0`: content `-> infinity`. This corresponds to a postulate that forbids everything but one outcome — maximally informative if true, and it is correct that such a postulate is nearly certain to die on first contact. The divergence is real and nameable.
FVP-4 FIXED POINTS	No interior fixed point. The invariant of interest is additive: two independent postulates over the same space have content `-ln(q_1 q_2) = H_1 + H_2`. **Content adds under independent conjunction**, which is the property that makes it a log measure and the reason to prefer this form over 1−q.
FVP-5 IDENTIFIABILITY	Poor, and stated as such. Two analysts enumerating outcomes differently will assign different q to the same postulate. **The measure is identifiable only relative to a published outcome list.** Publish the list or do not publish the number.
FVP-6 FALSIFICATION	The measure itself is not falsifiable; it is a bookkeeping definition. What is falsifiable is the *claim that it helps*: preregister a set of postulates with content scores, test them in a fixed budget, and check whether the high-content ones produced more parameter revision per unit cost. **If they did not, the heuristic is useless and should be dropped.**
FVP-7 REPRODUCIBILITY	Two independent analysts given the same postulate and the same outcome list should agree on q to within a stated tolerance. **This is checkable and it is the gate that decides whether the heuristic is usable at all.** Untested here; it is offered as a triage device with that limitation on the label.

37.11 Introspective report is data, not verdict

The sharpest version of the opening objection is not that inner states are unobservable but that the reports we substitute for them are unreliable. Nisbett and Wilson (1977) made the case directly: people confidently report the causes of their own judgements and choices, and those reports are frequently constructions rather than observations — plausible accounts assembled after the fact rather than readouts of process.

This manual accepts that result and it does not weaken the method. **It changes what a report is for.**

```
WRONG USE   The subject reports "I found the second one twice as loud",
       and we record that the second one was twice as loud in their experience.
            This treats the report as a verdict about an inner quantity.

RIGHT USE   The subject emits the number 2 under a specified procedure.
        We record the number as behaviour. We then ask whether the numbers
       emitted under that procedure have the structure the postulate predicts:
       consistency under repetition, transitivity, agreement across observers,
       the right functional form against a physical variable, and invariance
            when the procedure's irrelevant details are changed.
```

The right use never requires the report to be accurate about anything. It requires the *pattern of reports* to have properties that would be very unlikely if the postulated structure were absent. This is exactly how a thermometer is validated: nobody asks whether mercury knows the temperature. **A report is an instrument reading, and instrument readings are validated by their behaviour as a system, not by their sincerity.**

Nisbett and Wilson's finding constrains one class of report specifically: *causal* reports about one's own processing. It constrains magnitude reports far less, and Chapter 38 shows how well those pass the structural tests. **Reports are worthless as testimony and useful as readings, and the difference is whether you validate them by their structure.**

37.12 Stating your own subjective structure as a postulate

Here is the protocol, and it is meant to be executed, not admired. It is the personal-scope version of Chapter 24's survey.

Step A — Write the belief in its native, useless form. Whatever you actually say to yourself. "I'm sharpest in the morning." "I need conflict to think clearly." "Social events drain me." Write it exactly as you say it; the gap between this sentence and the one you finish with measures how much you did not know you were assuming.

Step B — Name the hidden quantity. In each of those sentences one unobservable does all the work: sharpness, clarity, drain. Name it, give it a symbol, and stop using the word.

Step C — State its structure. The postulate proper. It must answer four questions: what kind of quantity is it (a level? a rate? a capacity?); does it have a true zero; what makes it move; what observable does it touch. Badge it. Date it.

Step D — Derive at least three consequences, one of which must be uncomfortable. A consequence you would be glad to observe is worth little; you will find it whether or not it is there.

The uncomfortable consequence carries the content.

Step E — Write the forbiddance list. For each consequence, the specific observation that would contradict it, with numbers where numbers are available.

Step F — Choose the instrument and the sample size before observing. Then observe, then stop and compute. The stopping rule is part of the design; choosing when to stop after seeing the data is the commonest way a personal experiment becomes theatre.

Step G — Report to yourself in the same words you would use to report to a stranger. The gap between the two sentences is the size of the self-deception; record it in the ledger.

37.13 The freeze, and the skill score

Step F contains the move that separates a personal experiment from a personal story: **the freeze**. Paper E §5.3 requires it at institutional scale; it is more necessary at personal scale, because the person fitting the model and the person being measured are the same and there is no reviewer.

```
THE FREEZE

1. Split the record in two by TIME, never at random. Early half = training.
2. Fit the parameters on the training half only.
3. WRITE THE PARAMETERS DOWN. They are now frozen. No refitting.
4. Predict the held-out half with the frozen parameters.
5. Compare against four baselines. Report all five numbers, including the ones
   that embarrass you.
```

The four required baselines (Paper E §5.3) are **persistence** (tomorrow equals today), **direct-input** (output is a rescaling of input, no storage), **matched autoregressive** (same free-parameter count, no theory in it), and the **shuffled-input control** (the same frozen model, input series randomly permuted). The fourth is usually omitted and it is the decisive one: if your model predicts the held-out data just as well with the inputs shuffled, it is not using the inputs, and whatever it is doing has nothing to do with the structure you postulated.

The comparison statistic is the ordinary forecasting skill score:

> **HELD-OUT SKILL SCORE [STANDARD STATISTIC].**
>
> ```
> SS = 1 − RMSE_model / RMSE_baseline
> ```
>
> `SS = 0` means the model matched the baseline. `SS > 0` means it beat it. `SS < 0` means the baseline beat the model. Report one `SS` per baseline. **The headline number is the smallest of them**, not the largest.

```
FVP — HELD-OUT SKILL SCORE    SS = 1 − RMSE_model / RMSE_baseline
```

Gate	Verdict
FVP-1 DIMENSIONS	Both RMSEs are in the units of the predicted variable; the ratio is dimensionless and SS is dimensionless. Sound. Requires that both models predict *the same variable on the same points*; a skill score comparing predictions of different transforms of the outcome is meaningless.
FVP-2 SCALE TYPE	RMSE requires at least an interval scale on the outcome. **If the outcome is ordinal — a 1-to-5 rating — RMSE is not licensed and the honest substitutes are rank-based** (e.g. proportion of correctly ordered pairs). Chapter 53 is about this constraint; it bites here immediately and most personal analytics ignore it.
FVP-3 LIMITING CASES	`RMSE_model -> 0`: `SS -> 1`, perfect prediction; treat as a red flag for leakage rather than success. `RMSE_model -> RMSE_baseline`: `SS -> 0`, no information beyond the baseline. `RMSE_model >> RMSE_baseline`: `SS -> -infinity`, unbounded below, which correctly expresses that a model can be arbitrarily worse than doing nothing. `RMSE_baseline -> 0`: undefined, and correctly so — a baseline that is already perfect leaves nothing to explain.
FVP-4 FIXED POINTS	`SS = 0` is the invariant of interest: the locus where the theory contributes nothing. **The relevant fixed point for the shuffled-input control is also `SS = 0`, but with the opposite reading** — against a shuffled control, `SS = 0` means the model was never using its inputs and the theory is dead regardless of how well it fitted.
FVP-5 IDENTIFIABILITY	Fully identifiable given the frozen parameters and the held-out points. **Identifiability fails the instant any refitting occurs after the split**, and there is no statistical repair for that; the only defence is the written frozen parameter list.
FVP-6 FALSIFICATION	The falsification rule for a personal postulate: **the postulate fails if `SS <= 0` against any one of the four baselines.** Not the average. Any one. In particular, beating persistence while failing the shuffled control is a failure, and it is the most common failure.

Gate	Verdict
FVP-7 REPRODUCIBILITY	Recompute on a second, later, non-overlapping stretch of record — the personal equivalent of a second laboratory. **If the frozen parameters do not transfer to the later stretch, the postulate is at best indexed to a period and must be restated with that index in it.**

37.14 Why this is engineering and not epistemology

A reader may reasonably ask why a manual about building rooms and institutions has spent a chapter on philosophy of science. The answer is that in Volume Two the postulates are *load-bearing members*. Volume One could afford less rigour here because its quantities were externally observable; from Chapter 38 onward every design decision rests on a claim about an inner quantity, and if those claims are the wrong kind of object then every specification downstream is a specification for a building that will not stand. **The badge convention is a load rating. The FVP is the structural calculation. The freeze is the proof load test.**

There is one further reason, and it is the one that motivates the chapter's placement here rather than at the front of the book. **Postulates are stored order.** They sit in Q exactly as institutions do. They decay without maintenance (delta: you forget why you believed it), they convert experience into structure at some rate (eta: how readily a surprise updates you), they cost something to maintain (C: the effort of holding a belief you have not confirmed lately), and they occupy headroom (h: a mind full of unexamined postulates has no capacity for a new one). The Amendment Ledger of §37.9 is a delta-reduction device in exactly Chapter 3's sense. The freeze is an eta-raising device in exactly Chapter 4's sense. **The chapter is an application of the storage law to the contents of your own belief, and Chapter 40 will make that formal.**

The Specification

Numbered, checkable requirements. A postulate that fails any of these is not admitted to print, to a design document, or to a decision.

S37.1 — Badge in the same sentence. Every assertion that is not a theorem carries one of the six badges of §37.3, in the sentence in which it is first stated. No footnoted badges.

S37.2 — Date every postulate. A postulate without a date cannot be distinguished from a post hoc fit and must be treated as one.

S37.3 — Forbiddance list before observation. Write the list of observations that would kill the postulate, in operational terms, before the first observation is taken. The list is part of the postulate; a postulate cited without its list is cited incompletely.

S37.4 — Executable operational definitions. Every term in the forbiddance list must be applicable by a person who does not know the hypothesis, yielding the same classification. Test it by handing the definition to someone blind and checking agreement.

S37.5 — Asymmetry check. For each derived consequence, write the consequence the *negation* would produce. If they are the same, delete the consequence — it is not evidence.

S37.6 — Minimality audit at every revision. Delete any clause that forbids nothing; record the deletion. A postulate's word count should not grow monotonically across revisions.

S37.7 — Amendment ledger, maintained. One row per revision, forbiddance-added and forbiddance-removed columns filled. Three consecutive net-negative revisions triggers mandatory abandonment review.

S37.8 — Time-ordered split and freeze. Split by time, fit on the earlier portion, write the parameters down, do not refit.

S37.9 — Four baselines, all reported. Persistence, direct-input, matched autoregressive, shuffled-input. The headline is the minimum of the four skill scores.

S37.10 — Scope declaration. State whether the postulate is single-subject, group-indexed, or general. **A single-subject postulate may not be stated in the third person.** "I appear to recover analytic capacity in proportion to break length" is admissible. "People recover analytic capacity in proportion to break length" is a different postulate requiring FVP-7 evidence.

S37.11 — Scale-type declaration. State the Stevens scale type of every measured term (Chapter 53). If a formula divides two terms, both must be ratio-scaled with a true zero, or the formula is demoted to a heuristic ordering in print, visibly.

S37.12 — Publication of survival, not of truth. Report as "survived N rounds against the following forbiddance list", with the list attached. Never report a postulate as confirmed, established, or proven.

S37.13 — Decline path. Any protocol that asks another person to supply readings must state how they decline, and declining must cost them nothing (CANON §5.4). This binds personal-postulate work the moment it involves anyone else's reports.

The Failure Modes

FM-37.1 . THE UNBADGED POSTULATE

Damages: the reader's ability to calibrate. Downstream, `Q` — an unbadged postulate accumulates authority it never earned and becomes load-bearing without a load rating.

Signature: a statement about an unobservable, written in the indicative, with no test attached and no indication that one is owed. Often recognisable by a citation chain in which the earliest reference is a proposal and the latest is a premise.

Detection: take any paragraph and try to write a forbiddance list for each sentence. Sentences for which you cannot are either badged, badgeable, or contraband.

Fix: badge it retroactively and audit everything built on it. If the audit shows the structure above it cannot be supported by a [POSTULATE], the structure comes down, not the badge.

FM-37.2 . THE RETREATING POSTULATE

Damages: content. `H_forbid` shrinks monotonically toward zero while the postulate's apparent sophistication grows.

Signature: every disconfirmation is met with a qualification. The postulate has more clauses than it has ever had predictions. Nobody can state its current form in one sentence.

Detection: the amendment ledger. Net forbiddance negative across revisions.

Fix: abandon and restate from scratch. A postulate that has retreated three times has usually been replaced by a different postulate anyway; write the replacement down honestly and date it today, not at the original's date.

FM-37.3 . THE POSTULATE THAT FORBIDS NOTHING

Damages: the entire enterprise. It is not a weak postulate, it is not a postulate.

Signature: every possible outcome is described in advance as consistent. The proponent is delighted by all results.

Detection: the hostile-reviewer test (§37.8, Disguise 1), plus `q = 1` in Table 37.2.

Fix: demote to [METAPHOR] and stop citing it as support. A metaphor can be excellent. It cannot be evidence.

FM-37.4 . THE CIRCULAR INSTRUMENT

Damages: identifiability (FVP-5) and, silently, everything fitted with it.

Signature: the instrument that measures the outcome is derived from the same self-report that generated the postulate, so the postulate is being tested against a transformation of itself. In personal work this is nearly universal: you postulate that you are drained by X, then you measure drain by asking yourself how drained you feel after X.

Detection: Paper E §6.3's question, transposed to personal scope: **can you demonstrate that your outcome proxy is not a repackaging of your input?** If the outcome and the input come from the same sentence at the same moment, they are.

Fix: separate the channels — outcome by behavioural trace (work completed, error rate, timing), or by a different-modality report taken hours later and recorded without seeing the earlier one.

FM-37.5 . THE BORROWED CEILING

Damages: the reader's calibration, by importing authority across a boundary it does not cross.

Signature: "psychophysics established that subjective quantities are measurable, therefore my subjective quantity is measurable." Chapter 38 establishes something narrower: that *certain* subjective magnitudes, under *specific* procedures, yield ratio-scale numbers that replicate. **The existence proof licenses the attempt, not the result.**

Detection: look for the word "therefore" spanning two different quantities.

Fix: run the procedure on your quantity. An existence proof tells you the method is worth the cost; the cost is still yours to pay.

FM-37.6 . THE POST HOC FIT REPORTED AS A TEST

Damages: falsification (FVP-6). The reported evidence has no evidential value.

Signature: the model has as many free parameters as it has data points and fits beautifully. No held-out set. No frozen parameters. A story that accounts for every observed feature.

Detection: ask for the date of the parameter values and the date of the data. Ask for the held-out RMSE. Silence on either is diagnostic.

Fix: the freeze (§37.13). If the record is too short to split, say so and report the fit as exploratory, which is honest and useful; exploratory results are how good postulates are born.

FM-37.7 . THE IDENTITY POSTULATE

Damages: everything, and it is the hardest to fix, because the cost of falsification is not intellectual.

Signature: a postulate about oneself whose abandonment would require abandoning a self-image. "I am the kind of person who sees what others miss." The forbiddance list, if honestly written, contains observations the person cannot afford to make, so the list is quietly never written, or written in terms too vague to bind.

Detection: ask what changes about your life if the postulate is false. **If the answer includes who you are rather than what you do, the postulate is identity-loaded and its evidence must be discounted accordingly.**

Fix: two structural moves. Delegate the classification to someone blind to the hypothesis; and precommit the response, writing down in advance the change of practice that follows from the postulate's death, so death has a defined and survivable consequence. Chapter 39's framing register is the long-run treatment: an identity resting on one framing is an identity with `u = 0`.

FM-37.8 . THE CONSEQUENCE NOBODY RECORDED

Damages: the entire test, silently. This is the most frequent failure in practice and it looks like nothing at all.

Signature: the postulate was stated, the consequences derived, the observations made — and no artifact exists. The person "noticed" the results, and recall is then reconstructed under the influence of the postulate.

Detection: ask for the file. If the answer is a memory, there is no test.

Fix: any artifact — a text file, a paper notebook, a dated note. Chapter 4's arithmetic applies unmodified: an observation not recorded has `eta = 0` and contributes nothing to `Q`, however vivid it was.

FM-37.9 . THE MISSING SHUFFLED CONTROL

Damages: the causal claim specifically, while leaving the predictive claim intact and thereby disguising the damage.

Signature: a model that beats persistence is reported as evidence for the postulated mechanism. No shuffled-input comparison is reported.

Detection: look for the fourth baseline. Its absence is the signature.

Fix: run it. It costs one line of code and one permutation. See the Worked Example below, where it is the number that does the real work.

Worked Example — THE MORNING RESERVOIR

One operator, one postulate, carried from the useless native sentence to a frozen held-out test and an amendment. Every number below is computed. **The readings are a stipulated worked case chosen to demonstrate the arithmetic of the protocol; they are not empirical data about any person, and the chapter claims nothing about mornings.**

1. Step A — the belief in its native, useless form

> "I'm sharpest in the morning and I'm useless after lunch."

Two unobservables (*sharp*, *useless*), one implied cause (lunch), one implied structure, no quantities, no forbiddance. As stated it is compatible with every possible observation, so `q = 1` and `H_forbid = 0` — the starting condition of nearly every belief anyone holds about their own functioning.

2. Steps B and C — the named quantity and the postulate

The hidden quantity is a capacity for one class of work. Name it R, for reservoir. Restrict the class: **novel analytic problem-solving**, operationally *tasks requiring a solution method not previously used by this operator on this problem type*. Everything else is out of scope and is measured separately as a control.

> **THE MORNING RESERVOIR [POSTULATE, dated at the top of the record].** There is a single scalar capacity R for novel analytic work. R is depleted by time spent on that class of work and by nothing else. R recovers fully overnight and is **not** replenished by breaks within the day. Observed performance on a novel analytic task is monotone increasing in R at the moment the task is attempted.

Note the structure of the clauses. "By time on that class of work and by nothing else" forbids depletion by meetings, administrative work, emotional events. "Not replenished by breaks" forbids mid-day recovery. "Recovers fully overnight" forbids carry-over across days. Each clause does work; that is S37.6 satisfied at birth.

3. Step D — four consequences, one of them uncomfortable

```
D1  Within a working day, scores on the analytic class decline monotonically
```

```
     across successive blocks.

  D2  On days containing a long midday break, the block after the break falls
      ON the same decline curve. It does not rise.        <-- UNCOMFORTABLE

  D3  The control class (routine administrative work) shows no comparable decline.

  D4  A day that starts later shows the same decline SHAPE, shifted, not compressed:
      the decline tracks time-on-task, not clock time.
```

D2 is the uncomfortable one and it is where the content is: the operator takes breaks and would like them to be doing something. **A consequence you would be pleased to see falsified is exactly the consequence that carries information.**

4. Step E — the forbiddance list, written before observing

```
   F1  kills D1:  a non-monotone block sequence, in the mean over >= 20 days.
  F2  kills D2:  block 3 on break days exceeding the no-break decline prediction
                 by more than 0.5 pooled standard deviations.
  F3  kills D3:  the control class declining at more than one third the analytic rate.
   F4  kills D4:  late-start days showing a decline slope differing from early-
start
                 days by more than 25%.
```

Numbers, in advance, with tolerances. The postulate now forbids things.

5. Step F — the instrument and the design

Four 75-minute working blocks per day. At the end of each, the operator gives a magnitude estimate of the block's analytic output against a personal modulus of 100 assigned to a pre-selected reference block (the method is Chapter 38's; here it is used, not justified). Forty working days: twenty-two with no long break, eighteen with a 90-minute break between blocks 2 and 3.

TABLE 37.3 — BLOCK MEANS (stipulated worked case)

condition	block 1	block 2	block 3	block 4	days
no long break	100.0	88.0	74.0	63.0	22
(sd)	11.0	12.0	13.0	12.0	
90-min midday break	100.0	87.0	95.0	81.0	18
(sd)	10.0	11.0	12.0	13.0	
control: admin	100.0	98.0	97.0	96.0	22

6. Computation

Linear fits on the block index (computed):

```
no-break analytic:  slope = -12.5000 per block,  intercept 112.5000,  R^2 = 0.99808
break-day analytic:  slope =  -4.9000 per block,  intercept 103.0000,  R^2 = 0.56428
 control (admin):     slope =  -1.3000 per block
```

In the house register, log-ratios against block 1 (computed):

```
ln(88/100) = -0.1278     ln(74/100) = -0.3011     ln(63/100) = -0.4620

log-slope = -0.1559 nats per block,  R^2 = 0.99639
implied per-block retention (1 - delta) = exp(-0.1559) = 0.8556
implied per-block delta                 = 0.1444
```

That last line is the first substantive result. The no-break decline is very well described by a constant fractional loss per block — that is, by a leaky integrator with no input (Chapter 2). R behaves like Q under `Q_{t+1} = (1-delta)Q_t` with `delta = 0.1444` per 75-minute block. The postulate predicted only monotonicity, not an exponential. **The exponential form is a free finding, and free findings are how postulates get sharpened.**

The D2 test — the uncomfortable one:

```
frozen no-break decay applied forward from break-day block 1:
    predicted block 3 = 100.0 * exp(-0.1559 * 2) = 73.207

observed block 3 on break days                   = 95.000

difference                                       = 21.793
pooled sd (n1=22, n2=18)                         = 12.5625
Cohen d                                          =  1.7347
t                                                =  5.4583   (df = 38)
```

The tolerance in F2 was 0.5 pooled standard deviations. The observed excess is **1.73** pooled standard deviations, more than three times the tolerance, in the forbidden direction.

D2 is falsified. The Morning Reservoir postulate, as stated, is dead.

The other consequences:

```
D1  supported:  monotone decline in the no-break condition, R^2 = 0.99808
D3  supported:  control slope -1.3 vs analytic -12.5, a factor of 9.615;
          F3's threshold was one third, i.e. a factor of 3. Passed with margin.
D4  not tested: the operator did not vary start time. Recorded as untested,
                NOT as supported.
```

D4's row is the one to imitate. **An untested consequence is reported as untested, never folded into the supported column because the postulate "did well overall".**

7. The amendment, done legitimately

The clause "not replenished by breaks" is dead. The amendment must add forbiddance, not just remove it (§37.9), so the replacement clause is quantitative:

> **THE MORNING RESERVOIR, v2 [POSTULATE, dated at amendment].** As v1, except: R is partially replenished by breaks, and the replenishment is a saturating function of break length — specifically, recovery approaches a ceiling, so that doubling break length yields less than double the recovery.

Ledger row:

```
   version | clause changed             | forbiddance ADDED        | REMOVED | net
    --------+-----------------------------+---------------------------
+---------+-----
    v2   | breaks replenish, saturating  | ordering + concavity across | "no    | +
         |                              | three break lengths        | recovery"|
```

The new forbiddance is specific and testable on fresh data: with 20-, 45- and 90-minute breaks, recovery must be ordered `r(20) < r(45) < r(90)` **and** concave, `r(90) − r(45) < r(45) − r(20)`. That forbids a linear response and it forbids a step response. The net column is positive, so the amendment is progressive rather than a retreat.

8. The freeze and the four baselines

For the follow-up round the operator wants a predictive model, not just a contrast. The estimation model is Paper E §5.2's minimal form at personal scope:

```
    Y_{t+1} = a Y_t + b S_t + noise        a = 1 − delta,   b = eta
```

with Y the day's analytic output index and S the day's recovery input. The record is 40 days. Split by time at day 24. **The series used below is synthetic, generated by the stated rule `Y_{t+1} = 0.72 Y_t + 30 S_t + N(0, 2.0)` with a fixed seed, and it is here to demonstrate the arithmetic of the freeze, not to report a result about anyone.**

```
    FROZEN PARAMETERS, fitted on days 1–24 only:
        a = 0.7018      ->  delta = 0.2982
        b = 31.3083     ->  eta   = 31.3083

   (True generating values were a = 0.72, b = 30. Recovery is good, which is what
    one expects when the model form is correct and the noise is modest; it is not
      a result, it is a check that the fitting procedure works.)

    HELD-OUT PERFORMANCE, days 25–40, frozen parameters, computed:

      model                  RMSE =  1.5164    skill vs persistence =  0.8473
      persistence            RMSE =  9.9295    skill vs persistence =  0.0000
      direct-input           RMSE = 22.8039    skill vs persistence = −1.2966
     matched AR(1), c=0.9838  RMSE =  9.7719    skill vs persistence =  0.0159
     shuffled-input control   RMSE = 13.2567    skill vs persistence = −0.3351
```

Read the table in the order that matters. The model beats persistence by a wide margin, which by itself means little — persistence is weak on a series with real input variation. It beats the matched AR(1), which is the first result with content: **a one-parameter memory model with no theory in it scores 0.0159 against persistence, essentially nothing, so the improvement is not coming from autocorrelation alone.** The direct-input baseline is catastrophically bad, which tells us the output is not a rescaled copy of the input — there is genuine storage.

And then the fourth row, which is the one the whole protocol exists to produce. With the same frozen parameters and the input series permuted, held-out RMSE rises from 1.5164 to 13.2567 — a factor of 8.74 — and the skill score goes *negative*. The model is worse than doing nothing when its inputs are scrambled. **That is the result: performance depends on the specific temporal alignment of inputs with outputs, which is what a storage claim asserts and what a model**

that had merely learned the series' shape would not show. Had the shuffled RMSE come back near 1.5, the exercise would have been dead — and dead invisibly, because every other row looks excellent.

9. What was and was not established

```
    ESTABLISHED (as survival, not as truth):
    - a within-day monotone decline in the analytic class, well fitted by a constant
        fractional loss of 0.1444 per 75-minute block
      - the control class does not show it (factor of 9.6 in slope)
      - the frozen storage model beats all four baselines on held-out data,
        including and especially the shuffled-input control

    KILLED:
     - the clause "breaks do not replenish", at d = 1.73, three times the pre-
set tolerance

    UNTESTED, and recorded as untested:
      - D4, the time-on-task versus clock-time distinction
    - the saturating-recovery clause of v2, which is the next round's business

    NOT CLAIMED:
      - anything about mornings in general
    - anything about any other person          [scope: single subject, S37.10]
    - any account of WHY the decline occurs; the postulate is about structure,
        not mechanism, and no mechanism was tested
```

10. The rule, extracted

A postulate about your own interior is legitimate exactly when you have written down, before looking, the observation that would make you stop believing it — and then looked. Everything else in this chapter is machinery for making that one sentence executable.

Chapter 38 supplies the reason to bother. When this method was applied to subjective magnitude by Weber, Fechner and Stevens, it did not merely produce arguments. **It produced ratio-scale numbers, from human reports, that replicate across observers and laboratories.** The wall between the inside and the instruments is not solid. It has been crossed, repeatedly, by exactly the method this chapter specifies, and the crossing is a century old.

Chapter 38 . The Register Perception Actually Uses

Chapter 37 argued that an unobservable quantity can be studied by postulating a structure and testing the consequences. That argument is only as good as its best case, and this chapter is the best case. It is also the load-bearing evidence for the whole of Volume Two, so it is presented in full rather than in summary.

The claim to be established is narrow and it is enormous:

> **Subjective magnitude has been measured. Not ranked, not rated, not scored — measured, on a scale that supports ratios, from reports given by ordinary human beings under a procedure that takes about twenty minutes to run, and the measurements replicate.**

If that claim is false, Volume Two is a book of metaphors and should be read as one. If it is true, then the question about the inside is not *whether* it can be measured but *how well*, at what cost, and with what error bars — which is an engineering question, and this is an engineering manual.

The claim is true. It has been true since 1957, it rests on work begun in 1834, and the reason it is not more widely known is that the result belongs to a specialised subfield whose textbooks are not read outside it. This chapter reproduces the result, the arithmetic, and — with equal care — its limits, because the limits are severe and a reader who takes the result without them will misuse it within a page.

And there is a second result, which is the one that gives Volume Two its shape: the register that perception uses is logarithmic. That is not a metaphor either. It is the empirical content of Weber's law, it is the form of Fechner's law, it is what taking logarithms of Stevens' law reveals, and it is why every quantity in this framework — surplus, pitch, sensation, moment — is a logarithm (CANON_INTERNAL §1). The formalism is logarithmic because the lens is.

The Law

38.1 Three laws, three centuries, one register

```
 WEBER    1834    dI / I = c                    the JND is a constant RATIO
FECHNER  1860   S = k ln(I / I_0)          sensation is LOGARITHMIC in intensity
STEVENS  1957   psi = k I^a               sensation is a POWER FUNCTION of intensity

 and, taking logarithms of the third:

                    log psi = log k + a log I      LINEAR in log–log
```

Three statements, each an advance on the last, and the third did not abolish the register — it relocated it. Weber's law is a statement about ratios. Fechner's law is a logarithm outright. Stevens' law is a straight line *in log-log coordinates*, which is to say a statement that equal ratios of stimulus produce equal ratios of sensation. **All three say that perception operates on ratios and not on differences. They disagree only about what the ratio does when it arrives.**

That is the register claim, and it is the one Chapter 36's Lens Theorem is built on. It survives the Fechner-to-Stevens transition intact, which matters, because the transition killed Fechner's

specific functional form and a reader could be forgiven for thinking it killed the logarithm with it. It did not. It replaced *log in, linear out* with *log in, log out*.

38.2 Weber (1834): the just-noticeable difference is a ratio

Ernst Heinrich Weber, working on touch and on lifted weights, reported that the smallest detectable increase in a stimulus is proportional to the stimulus already present.

```
WEBER'S LAW   [EMPIRICAL REGULARITY]

  dI / I = c        c = the Weber fraction, a constant for a given modality
                            over the mid-range of that modality
```

Two properties make this a strong result rather than an observation. First, it is about *discrimination performance*, which is behaviour: whether an observer can tell two stimuli apart is settled by a forced choice and a tally, with no inner quantity invoked anywhere. Second, it holds across modality after modality over the middle of each modality's range, and it fails at both ends in a characteristic and reproducible way — near absolute threshold `c` rises, near saturation it rises again. **A law that fails at the edges in the same way everywhere is a law with a mechanism behind it, not a coincidence.**

`FVP — WEBER'S LAW    dI / I = c`

Gate	Verdict
FVP-1 DIMENSIONS	`dI` and `I` share the physical units of the stimulus, so `c` is dimensionless. Sound, and the dimensionlessness of `c` is the substantive content: it makes the fraction comparable *across* modalities, which is what later allows Weber fractions and Stevens exponents to be related at all.
FVP-2 SCALE TYPE	`I` must be ratio-scaled with a true zero — physical intensity, luminance, force, concentration all qualify. **The law is not statable for a stimulus dimension that lacks a true zero**, and this is a real restriction: for pitch, for hue, for position, Weber's law in this form does not apply and the analogous statement must be constructed differently.

Gate	Verdict
FVP-3 LIMITING CASES	`I -> 0`: `dI -> 0`, predicting infinitely fine discrimination near zero. **This is false and known to be false**; absolute threshold intervenes, and the standard repair is `dI/(I + I_r) = c` with `I_r` a small resting term. The divergence corresponds to a real, named phenomenon (absolute threshold), which is what FVP-3 requires. `I -> infinity`: `dI` grows without bound, which fails at saturation and, in pain-adjacent modalities, at tissue damage.
FVP-4 FIXED POINTS	The invariant is the fraction itself: `c` is unchanged under `I -> aI` for any positive `a`. **Scale invariance of discrimination is Weber's law's entire content**, and it is the property that forces every downstream law into a logarithmic register.
FVP-5 IDENTIFIABILITY	`c` is recoverable from forced-choice discrimination data by standard psychometric-function fitting; a JND is conventionally the 75%-correct point in a two-alternative task. **The convention must be stated with the number**, because a `c` defined at 75% and a `c` defined at 84% are different numbers for the same observer.
FVP-6 FALSIFICATION	Killed by a modality showing a constant *absolute* JND over its mid-range, or by `c` varying systematically with `I` in the mid-range beyond measurement error. Both outcomes occur at the edges of ranges and neither occurs in the middle; the law is therefore alive as a mid-range law and dead as a universal one, and that is exactly how it is reported here.
FVP-7 REPRODUCIBILITY	**This gate is passed and it is the reason the chapter exists.** Weber fractions recover across observers and laboratories to within modality-typical scatter, using apparatus available in 1834 and apparatus available now. Individual differences are real and are themselves stable.

38.3 Fechner (1860): the logarithm

Chapter 37 gave the derivation and the FVP table in full, as the exemplar of the Postulate Method, and they are not repeated. The one-line reminder:

```
postulate:  all JNDs are subjectively equal     [POSTULATE, Fechner 1860]
+ Weber:    one JND = c I
integrate:  S = k ln(I / I_0)                     FECHNER'S LAW
```

What belongs in this chapter is a consequence Chapter 37 stated but did not develop. If sensation is counted in JNDs, then a modality's whole dynamic range contains a **finite, computable number of discriminable steps**:

> **THE DISCRIMINABLE-STEP COUNT [DERIVED from Weber].** For a modality with Weber fraction `c` over a dynamic range `R = I_max / I_min`,
>
> `N_jnd = ln R / ln(1 + c)`

TABLE 38.1 — DISCRIMINABLE STEPS ACROSS A DYNAMIC RANGE (computed)

c	R = 10	R = 100	R = 10^4	R = 10^6	ln N_jnd at R = 10^6
0.01	231.4	462.8	925.6	1388.4	7.2359
0.02	116.3	232.6	465.1	697.7	6.5477
0.05	47.2	94.4	188.8	283.2	5.6460
0.10	24.2	48.3	96.6	145.0	4.9764
0.20	12.6	25.3	50.5	75.8	4.3278
0.33	8.1	16.1	32.3	48.4	3.8804

The last column is the one that connects this chapter to the rest of the volume. `N_jnd` is a count of distinguishable categories along one perceptual dimension, and `ln N_jnd` is therefore the surplus ceiling for that dimension (Paper D, Cor 2.2; CANON_INTERNAL §1). A modality with a Weber fraction of 0.05 and six decades of range affords at most 5.646 nats of surplus, no matter what is done with it. A modality with `c = 0.33` affords 3.880. **The maximum possible moment along a perceptual dimension is fixed by that dimension's discrimination acuity and its dynamic range, and both of those are measurable with a forced-choice task.** This is Ashby's requisite variety arriving in first-person experience, and it arrives with a number attached.

`FVP — DISCRIMINABLE-STEP COUNT    N_jnd = ln R / ln(1 + c)`

Gate	Verdict
FVP-1 DIMENSIONS	`R` is a ratio of like units, hence dimensionless; `c` is dimensionless; `N_jnd` is a pure count. Sound.
FVP-2 SCALE TYPE	Requires `I` ratio-scaled (inherited from Weber). The output is a count, which is a ratio scale with a true zero, so ratios of step counts across modalities are licensed.

Gate	Verdict
FVP-3 LIMITING CASES	`c -> 0`: `N_jnd -> infinity`, infinitely fine discrimination gives infinitely many categories — correct in the model and correctly impossible in fact. `c -> infinity`: `N_jnd -> 0`, an observer who cannot discriminate has one category, which is right. `R -> 1`: `N_jnd -> 0`, no range means no steps. All three divergences name real phenomena.
FVP-4 FIXED POINTS	`N_jnd = 1` at `R = 1 + c`: the range that contains exactly one step. This is the definition of the JND recovered as a fixed point, which is a consistency check on the formula rather than a discovery.
FVP-5 IDENTIFIABILITY	Both parameters come from independent measurements — `c` from discrimination, `R` from threshold and saturation determination. **Fully identifiable, and this is unusual enough in the human sciences to be worth marking.**
FVP-6 FALSIFICATION	Killed by a direct count of discriminable steps (an identification task with feedback, counting categories reliably distinguished) disagreeing with the computed `N_jnd` by more than measurement error. **This is a real and frequently failed test**: absolute-identification experiments typically recover far fewer categories than the JND count predicts, because identification without a comparison stimulus is much harder than discrimination with one. **The formula therefore overstates usable categories and must be reported as an upper bound.**
FVP-7 REPRODUCIBILITY	Inherits Weber's FVP-7 for `c`; `R` is more variable across observers, particularly at the upper end where tolerance rather than physiology sets the limit. Report the range convention with the number.

38.4 Stevens (1957): the power law, and how it was obtained

Fechner's law dominated for nearly a century. S. S. Stevens broke it, and the break is instructive because it came not from a better argument but from **a better instrument**.

Fechner had reached sensation indirectly, by counting discriminations and postulating that the counted units were equal. Stevens asked observers to do something no one had thought worth trusting: **to assign numbers directly to sensations, in proportion.** Present a stimulus, call it 10. Present another. If it seems twice as strong, say 20. If it seems a third as strong, say 3.3.

This is the **method of magnitude estimation**, and the field's prior assumption was that it would produce noise.

It did not. It produced, for modality after modality, a relation of the form

```
STEVENS' POWER LAW   [EMPIRICAL REGULARITY]

    psi = k I^a

and equivalently, taking logarithms:

   log psi = log k + a log I          a straight line in log–log coordinates
```

with a stable for a given modality and procedure, and with the log-log fit typically excellent. Stevens set out the case in *On the psychophysical law* (Psychological Review, 1957) and, more combatively, in *To honor Fechner and repeal his law* (Science, 1961).

The exponent a is a measured conversion key between the physical register and the perceptual one (CANON_INTERNAL §3), and it is the single most useful number in this chapter.

`FVP — STEVENS' POWER LAW    psi = k I^a`

Gate	Verdict
FVP-1 DIMENSIONS	`I` carries physical units; `I^a` for non-integer a is dimensionally improper unless `I` is first expressed as a ratio to a reference. **The law must therefore be written `psi = k (I/I_ref)^a`, and the omission of `I_ref` in casual statements of it is a real dimensional error that changes k and hides the reference.** `psi` inherits the units of k, which are the units of the response modulus. State `I_ref` and the modulus or the constant is meaningless.
FVP-2 SCALE TYPE	`I` ratio-scaled. **`psi` is claimed to be ratio-scaled, and that claim is the whole point of the method and the whole locus of the controversy.** The claim is that observers can execute the instruction "make the number proportional to the sensation" — that they perform a ratio judgement, not a rank ordering with numerals attached. The evidence for it is internal-consistency evidence (§38.13), not direct. Chapter 53 examines it as the crux of the volume.

Gate	Verdict
FVP-3 LIMITING CASES	`I –> 0`: `psi –> 0` for `a > 0`, appropriate, though absolute threshold again truncates the domain; the standard repair is `psi = k (I – I_0)^a`. `I –> infinity`: `psi –> infinity`, unbounded, wrong at saturation, so the law is again a mid-range law. `a –> 0`: `psi –> k`, a modality with no sensitivity to intensity — a real limiting case, approached by some quality dimensions.
FVP-4 FIXED POINTS	The scale-invariance property is the invariant: `psi(bI)/psi(I) = b^a` for all `I`. **Equal stimulus ratios give equal sensation ratios everywhere on the scale, and the constant of proportionality is `b^a`.** This is the exact analogue of Fechner's `S(aI) – S(I) = k ln a`, with a ratio where Fechner had a difference, and it is what makes the exponent a conversion key.
FVP-5 IDENTIFIABILITY	a is identifiable from the log-log slope and is robust; k is not separately identifiable without fixing the modulus, since any rescaling of the observer's number use rescales k. **In practice a is the reportable parameter and k is a procedural artifact.** This asymmetry is fortunate: the exponent is the physically meaningful quantity and it is the identifiable one.
FVP-6 FALSIFICATION	Killed by systematic curvature in log-log coordinates beyond measurement error; by exponents that do not survive a change of modulus or range; or by cross-modality matches whose exponents disagree with the ratio predicted from the separate power functions (§38.13). **All three tests have been run. The first largely passes, the second largely fails, and the third largely passes.** That mixed verdict is the honest state of the result and it is stated as such below.

Gate	Verdict
FVP-7 REPRODUCIBILITY	**Passes at the level of ordering and approximate magnitude, fails at the level of the third decimal place.** Exponents replicate across observers and laboratories in rank order and to within roughly ten to twenty per cent for well-studied modalities; they shift systematically with stimulus range, instruction wording, and modulus. Report them as constants of a procedure, not of nature.

38.5 The existence proof, stated exactly

Here is what the preceding sections license and nothing more.

> **THE EXISTENCE PROOF.** For a class of subjective magnitudes — perceived loudness, brightness, heaviness, warmth, length, taste intensity, and others — there exists a procedure, executable in under an hour with ordinary human participants, that yields numbers with the following properties: they are ordered consistently with the stimulus; they are stable under repetition within an observer; they agree across observers in form and approximately in exponent; they support ratio operations in the sense that a stimulus ratio produces a predictable response ratio; and they cohere across modalities in a way that a scale of arbitrary numerals would not.

Four things this does **not** license, stated with the same emphasis:

1. **It does not license the conclusion that all subjective quantities can be measured this way.** Loudness has a physical correlate that can be varied precisely. Grief does not. The proof establishes that the wall between inner magnitude and instrumentation is not solid; it says nothing about where else it can be crossed. (Chapter 37, FM-37.5, the Borrowed Ceiling.)
2. **It does not license treating rating scales as ratio scales.** A one-to-five satisfaction item is not magnitude estimation and produces ordinal data. Chapter 53.
3. **It does not license the exponents as constants of nature.** §38.15.
4. **It does not settle whether the number an observer emits is "the sensation".** It settles that the numbers behave as a measurement system behaves. Chapter 37 §37.11 is the operative distinction: readings, not testimony.

38.6 The honesty clauses for this chapter

1. Every exponent below is **nominal, modality- and method-dependent**, reproduced from the published psychophysical literature as summarised in CANON_INTERNAL §3. **They vary with method, range and instruction and must be reported as measured constants of a procedure, not constants of nature.**
2. The ratio-scale status of magnitude estimation is **defended, not proved**. The measurement-theoretic literature (the *Foundations of Measurement* programme) has developed axiomatic conditions under which such judgements would constitute ratio measurement; whether human observers satisfy them is an empirical question that is still open.

3. Nothing here is a claim about consciousness, physics, or thermodynamics (CANON §5.1).
4. Where this chapter computes an illustrative table from stated parameters, the table is arithmetic, not data, and is labelled so.

The Engineering

38.7 The exponent table, read properly

TABLE 38.2 — PUBLISHED STEVENS EXPONENTS (nominal; CANON_INTERNAL §3)
with computed consequences

modality	a	class	I ratio for psi to x2	psi ratio for I x2	psi ratio for I x10
Loudness (1 kHz, SPL)	0.67	compressive	2.8138	1.5911	4.6774
Brightness	0.33	compressive	8.1698	1.2570	2.1380
Visual length	1.00	linear	2.0000	2.0000	10.0000
Taste (salt)	1.30	expansive	1.7044	2.4623	19.9526
Heaviness (lifted)	1.45	expansive	1.6129	2.7321	28.1838
Warmth on arm	1.60	expansive	1.5422	3.0314	39.8107
Electric shock	3.50	expansive	1.2190	11.3137	3162.2777

a < 1 compressive, a = 1 linear, a > 1 expansive. The three-way split is not a taxonomic convenience; it is a functional distinction with a plain reading.

Compressive modalities are the ones that must handle enormous physical ranges. Sound pressure across the audible range spans six orders of magnitude; luminance across a day spans more than ten. A compressive exponent packs that range into a manageable perceptual span at the cost of resolution: at a = 0.33, the world must get **8.17 times** more intense before the sensation doubles. This is a gain control, and the price of the gain control is that you cannot perceive the size of what you are looking at.

Expansive modalities are the ones that must trigger action. Warmth on the arm, at a = 1.6, triples in felt magnitude when the stimulus doubles. Electric shock, at a = 3.5, multiplies by **11.31** when the stimulus doubles and by **3162** when the stimulus goes up a factor of ten. **A system that must produce withdrawal before damage cannot afford compression, and the numbers show it has none.** Most senses compress; pain expands. That sentence is the shortest true summary of Table 38.2 and it has a clear functional logic behind it, though the logic is an interpretation and the numbers are the data.

Visual length at a = 1.0 is the linear case and it is the one to keep in mind whenever you design a display. Perceived length is very nearly proportional to actual length, which is precisely why bar charts work and why area and volume encodings — which pass through the compressive exponents for area and volume — do not.

38.8 What an exponent costs you

Column four of Table 38.2 is the design column. **It answers the question every engineer actually has: how much more of the physical thing must I supply for the recipient to notice a doubling?**

TABLE 38.3 — WORLD COST OF ONE PERCEIVED DOUBLING (computed)

modality	a	log-ratio required (nats)	in decades
Electric shock	3.50	0.1980	0.0860
Warmth on arm	1.60	0.4332	0.1881
Heaviness (lifted)	1.45	0.4780	0.2076
Taste (salt)	1.30	0.5332	0.2316
Visual length	1.00	0.6931	0.3010
Loudness (1 kHz, SPL)	0.67	1.0345	0.4493
Brightness	0.33	2.1004	0.9122

computed as ln(2)/a nats; a perceived doubling costs ln 2 = 0.693147 nats of response log-magnitude, and the exponent converts that into stimulus log-units.

Read the brightness row against the shock row. **To double perceived brightness you must supply 0.912 decades — a factor of 8.17 — of luminance. To double perceived shock you must supply 0.086 decades, a factor of 1.22.** The same subjective step costs a physical factor of 8.17 in one register and 1.22 in the other — or, measured in the log-units where the comparison is proper, 2.1004 nats against 0.1980, a ratio of 10.6. This is why "make it brighter" and "make it hurt less" are engineering requests of entirely different magnitudes, and why a budget that treats them as comparable will be wrong by most of an order of magnitude.

This table is Chapter 36's Lens Theorem in its operational form, with modality-specific keys attached. **Design in ratios; look up the exponent; convert.**

38.9 The loudness cross-check, and why FVP-1 was not pedantry

The loudness exponent is quoted as 0.67 and this is the row most often misused, because 0.67 is an exponent on **sound pressure**, not on sound intensity, and intensity goes as pressure squared. The FVP-1 gate on Stevens' law flagged exactly this, and here is what it costs to ignore it.

```
from Table 38.2:  pressure ratio for perceived doubling = 2.8138
intensity ratio  = 2.8138^2                               = 7.9177
in decibels      = 10 log10(7.9177)                       = 8.986 dB

implied intensity-domain exponent = 0.67 / 2              = 0.335
```

Now compare with the independent convention from acoustics. The sone scale defines a doubling of loudness, and the standard rule of thumb is that **a 10 dB increase doubles perceived loudness**. That rule implies an intensity-domain exponent of `ln 2 / ln 10 = 0.3010`, hence a pressure-domain exponent of **0.6021**.

```
from Stevens' 0.67 (pressure)   ->  perceived doubling at  8.99 dB
from the sone convention        ->  perceived doubling at 10.00 dB
discrepancy                                                1.01 dB
```

Two methods developed for different purposes, one from magnitude estimation and one from loudness-matching standards, agree on the location of a perceived doubling to within about one decibel. That is the kind of convergence that makes a measurement claim credible, and it is offered here precisely because it is *not* exact. The one-decibel gap is the method dependence of §38.15 showing up in the third significant figure, exactly where the honesty clause says to expect it.

38.10 Why the dispute lasted a century

A logarithm and a compressive power function look alike over a limited range, and this is computable rather than impressionistic. Fit a logarithmic law by least squares to a power function with `a = 0.33` over two decades of stimulus:

```
over I in [1, 100]  (two decades, 200 log-spaced points)

    best fit:  S = 0.6306 + 0.7475 ln I
    R^2      = 0.963165
    maximum relative deviation = 36.94 %

over I in [1, 10000]  (four decades)

    R^2      = 0.872523
    maximum relative deviation = 314.06 %
```

Over two decades a logarithm captures 96.3% of the variance of a power function. With nineteenth-century instrumentation, observer variability of ten to twenty per cent, and stimulus ranges rarely exceeding two decades, that is not distinguishable. Over four decades it becomes unmistakable: the maximum relative deviation goes from 37% to 314%.

Stevens did not win by argument. He won by extending the range and by asking a question whose answer distinguished the two hypotheses. That is the general lesson and it is worth extracting as a rule: *when two models fit your data equally well, the fix is not more data in the same range; it is data in a range where they disagree.* Chapter 37's content measure (§37.10) is the formal version of the same advice.

The second and decisive blow was qualitative. **A logarithmic law cannot produce an expansive relation at all** — `ln` is concave everywhere, so no logarithmic sensation scale can grow faster than its stimulus. Heaviness at 1.45, warmth at 1.6 and shock at 3.5 are not near-misses for Fechner's law; they are outside its range of possible outputs. One expansive modality, measured reliably, kills the law as a general account. Several were.

38.11 The step count and the ceiling on a moment

Return to Table 38.1 with the framework's question. `ln N_jnd` is the maximum surplus available along a single perceptual dimension, and Volume Two's forced consequence (CANON_INTERNAL

§1) is that **the maximum possible moment is `ln N`** — the log of how many kinds of thing the perceiver can distinguish.

Table 38.1 makes that ceiling concrete for one dimension at a time: between 3.88 and 7.24 nats across the parameter range shown. But a moment is not built from one dimension. If an observer attends to `m` perceptual dimensions whose discriminations are independent, the joint category count multiplies and the ceiling adds:

```
ln N_total = sum over dimensions of ln N_jnd(i)     IF the dimensions are independent
```

and that conditional is doing all the work. **Independence across perceptual dimensions is exactly the `u = 1` condition of Paper A, transposed to one observer's own channels, and it is not generally satisfied.** Loudness and brightness in a room co-vary with how many people are in it. The honest form is Paper A's:

```
N_eff = 1 + (m − 1) u       surplus = ln N_eff       [POSTULATE at this scope]
```

with u the mean non-parallelism between the channels. **Two perceptual channels that always move together are one channel, and the ceiling they set jointly is the ceiling one of them sets alone.** This is Chapter 39's subject and Chapter 39 derives it properly; it is flagged here because Table 38.1 is where the numbers to feed it come from.

38.12 Running a magnitude estimation

The procedure, in the form that survives the criticisms of §38.15 best. This is the instrument the rest of Volume Two borrows, so it is specified rather than described.

1. FIX THE CONTINUUM. One dimension, varied alone, with everything else held constant. If two things vary, you are not measuring either.

2. CHOOSE THE STIMULUS SET. At least 7 levels, GEOMETRICALLY spaced, spanning at least 1.5 decades. Geometric spacing, not linear: the register is logarithmic, so linear spacing wastes most of your points at the top.

3. RANDOMISE ORDER, fresh per observer. Sequential dependencies are real (§38.15).

4. INSTRUCT IN RATIOS, not in categories:
   "Assign any positive number to the first stimulus. For each one after, assign a number in PROPORTION to how strong it seems. If one seems half as strong, use half the number. Do not restrict yourself to whole numbers. Do not restrict yourself to any range."
   The instruction to use ratios and the removal of range restriction are both load-bearing. An instruction that mentions a scale from 1 to 10 destroys the ratio property and yields ordinal data.

5. FREE MODULUS, or a modulus in the MIDDLE of the range. A modulus at an extreme biases the exponent, because observers are reluctant to go far past an anchor.

6. REPLICATE each level at least twice per observer, separated in the sequence.

7. AGGREGATE BY GEOMETRIC MEAN, never the arithmetic mean. Responses are on a ratio

scale, so the arithmetic mean is the wrong summary and is dominated by whichever observer used the largest numbers.

8. FIT IN LOG–LOG. Regress log(response) on log(stimulus). Slope = a. Report R^2, the residual pattern, the stimulus range, the modulus policy, and the exact instruction text. WITHOUT THOSE FOUR, THE EXPONENT IS NOT INTERPRETABLE.

9. TEST THE RATIO PROPERTY DIRECTLY: pick stimulus pairs with equal physical ratios at different absolute levels; the response ratios must match. This is a check the observer could fail, and it is the one that earns the scale type.

Step 9 deserves emphasis. **The ratio-scale claim is not established by the fit; it is established by the invariance.** A power function fitted to numerals would still fit. What a merely ordinal response would not do is produce the *same* response ratio for the same stimulus ratio at different absolute levels. That test is cheap and it is skipped almost universally.

38.13 Cross-modality matching

Stevens' strongest internal-consistency argument came from a procedure that removes numbers altogether. Instead of asking an observer to say a number, ask them to **adjust one modality to match another**: squeeze a handgrip until the force feels as strong as the tone sounds loud.

The prediction is immediate and quantitative. If modality 1 has exponent a_1 and modality 2 has exponent a_2 against their respective physical continua, then matching 1 to 2 should produce a relation whose log-log slope is a_1 / a_2.

TABLE 38.4 — CROSS-MODALITY MATCHING PREDICTIONS (computed from Table 38.2)

match	to	predicted matching exponent a_1/a_2
heaviness	loudness	2.1642
warmth	brightness	4.8485
shock	length	3.5000
salt	loudness	1.9403

These are risky predictions in Chapter 37's sense: they are derived from measurements made with one procedure and tested with a completely different one that involves no numbers at all. Nothing in the matching task lets the observer consult the exponents. Stevens reported that matching exponents came out close to the predicted ratios across many modality pairs, and this result is the strongest evidence available that magnitude estimation is measuring something rather than recording a habit of number use.

It is also, honestly, evidence of *consistency* rather than of *correctness*. A systematic bias present in both tasks would produce the same agreement. The measurement-theoretic literature has pressed exactly this point. **The correct summary is: the numbers cohere across procedures that share no method, which is much more than a rating scale can claim and less than a physical measurement can claim.**

38.14 The exponent as a conversion key

CANON_INTERNAL §1 requires that every quantity in this framework be inter-convertible by exact constants, on the grounds that all of them are logarithmic. The exponent is the key that connects the physical register to the perceptual one, and now it can be stated precisely.

```
perceptual log-magnitude:   ln psi = ln k + a ln I

therefore:                  d(ln psi) = a * d(ln I)

ONE NAT of stimulus log-ratio  ->  a NATS of perceptual log-ratio.
```

The exponent is a linear map between two logarithmic registers, and that is exactly what a conversion key is (CANON_INTERNAL §1). Composing it with the volume's other keys:

```
  stimulus log-ratio (nats)  --  x a  -->  perceptual log-ratio (nats)
perceptual log-ratio     -- x 2.885390 --> octaves      (CANON_SONIC §4)
  perceptual log-ratio       --  x 1.442695  -->  bits
  perceptual log-ratio       --  x 3462.4681 -->  cents
```

Worked instance: a tenfold increase in luminance is `ln 10 = 2.302585` nats of stimulus log-ratio. At `a = 0.33` that is `0.759853` nats of perceptual log-ratio, which is `2.192` octaves in the volume's pitch register and `1.096` bits. A tenfold increase in shock intensity is the same 2.302585 nats of stimulus and `8.059` nats of perceptual log-ratio — `23.25` octaves, `11.63` bits. **The same physical decade lands as 1.1 bits in one register and 11.6 bits in another.** The conversion is exact given the exponent; the exponent is the uncertain part, and it is uncertain in the way §38.15 describes.

38.15 The honest limits, in detail

This section is not a hedge appended to a result. **It is part of the result**, and a reader who takes Table 38.2 without it will overreach within a page.

The exponent depends on stimulus range. Compressing the range of presented stimuli tends to raise a compressive exponent and lower an expansive one; observers spread their numbers over roughly the same span regardless of how much physical range you gave them. Teghtsoonian (1971) argued that this is systematic — that the subjective range an observer uses is approximately constant across modalities, which would make the exponent partly a consequence of dynamic range rather than a property of the sense. **If that is right, the exponent is a joint property of the sense and the experiment, and reporting it without the range is reporting half a number.**

The exponent depends on the modulus and the anchor. A modulus placed at one end of the range pulls the fitted exponent toward the anchor. Free-modulus designs and mid-range moduli reduce but do not eliminate this.

The regression effect. Observers under-use the extremes of whatever they are asked to adjust. Magnitude estimation (adjust the number to the stimulus) and magnitude production (adjust the stimulus to a given number) yield systematically different exponents, bracketing the true value. The standard repair is to run both and take the geometric mean of the two exponents; a single-direction study should say which direction it ran.

Sequential dependencies. A response is influenced by the preceding stimulus and the preceding response. This is why Step 3 randomises and why replicates must be separated.

Number-use habits. Observers prefer round numbers, powers of ten, and small integers. Poulton (1968) catalogued the biases in magnitude estimation at length and the catalogue has not been refuted; it has been managed by better designs.

Individual differences are large and real. Exponents vary across observers by a substantial fraction of their value. This is not only noise — the individual exponents are themselves reasonably stable on retest, which means the population figure is a central tendency of a genuine distribution, not an estimate of a single constant.

```
THE RULE THAT FOLLOWS, and it is binding on every use of an exponent in this volume:

 REPORT AN EXPONENT AS  a = value  (method, range, modulus, instruction, n)

An exponent quoted bare is a number without a referent. The values in Table 38.2 are
reproduced as NOMINAL published figures for orientation and comparison, and they are
not to be used as calibration constants for any specific application without
  re-measurement under the application's own conditions.
```

None of this touches the existence proof. That subjective magnitude can be measured on a ratio-supporting scale is established by the structure of the results — the log-log linearity, the within-observer stability, the cross-modality coherence, the cross-method convergence with the decibel convention. **What the limits establish is that the measurement has an error budget, like every measurement, and that the budget is larger than physics is used to and smaller than the human sciences usually admit.** That is the state of the art, and it is a good deal better than the state of the objection Chapter 37 opened with.

The Specification

S38.1 — Ratio instruction, verbatim. Any protocol claiming ratio-scale subjective data must publish its exact instruction text. Instructions that mention a bounded range, named categories, or integer-only responses do not yield ratio data and must be reported as ordinal.

S38.2 — Geometric spacing, minimum 7 levels, minimum 1.5 decades. State the range.

S38.3 — Geometric mean aggregation. Arithmetic means of magnitude estimates are prohibited.

S38.4 — Log-log fit with residual inspection. Report slope, intercept, R^2, and the residual pattern. Systematic curvature is reported, not smoothed.

S38.5 — Both directions, or declare one. Run estimation and production and report the geometric mean of the exponents; or run one and state which, with the direction of the expected bias.

S38.6 — Ratio-invariance test. Include equal-physical-ratio pairs at two absolute levels and report whether the response ratios matched. **A protocol without this test may not claim ratio scale.**

S38.7 — Full exponent citation format. `a = value (method, range, modulus, instruction, n)`.

S38.8 — Units before exponents. State the physical quantity the exponent is taken on (pressure vs intensity; luminance vs illuminance; concentration vs mass). Pressure-versus-intensity alone is a factor-of-two error in the exponent.

S38.9 — Upper-bound labelling of step counts. `N_jnd` is reported as an upper bound on usable categories, with the identification-task caveat attached.

S38.10 — Re-measure before calibrating. Published exponents orient a design; they do not calibrate one. Any application that depends on an exponent's value re-measures it under its own conditions.

S38.11 — Decline path. Participants may stop at any point without giving a reason and without cost, and this is stated to them before the first trial (CANON §5.4).

The Failure Modes

FM-38.1 . THE RATING SCALE WEARING A RATIO SCALE'S CLOTHES

Damages: FVP-2, and therefore every division and multiplication downstream.

Signature: a one-to-seven Likert item whose means are divided, multiplied, or compared as ratios. "Satisfaction rose 20%" computed from a bounded ordinal scale.

Detection: look for a bound. A response scale with a ceiling is ordinal at best; the ceiling itself proves the scale is not ratio, because a ratio scale has no maximum.

Fix: magnitude estimation (§38.12) if the quantity supports it; otherwise demote every formula containing the quantity to a heuristic ordering, visibly, per Chapter 53.

FM-38.2 . THE EXPONENT QUOTED BARE

Damages: reproducibility. The number cannot be checked, compared, or reused.

Signature: "the exponent for X is 1.4", with no range, modulus, instruction, or direction.

Detection: try to reproduce it. If the protocol is underspecified, it was quoted bare.

Fix: S38.7 format, or drop the number.

FM-38.3 . THE PRESSURE/INTENSITY CONFUSION

Damages: dimensions (FVP-1) by a factor of exactly two in the exponent, which is a factor of four in every conversion.

Signature: loudness computations mixing an exponent of 0.67 with intensity data, or 0.33 with pressure data. Results off by a factor of two in log-space.

Detection: the §38.9 cross-check. If your perceived doubling does not land within a decibel or so of 10 dB, you have the wrong domain.

Fix: state the domain before the exponent; convert with `a_intensity = a_pressure / 2`.

FM-38.4 . THE LINEAR PROMISE IN A LOGARITHMIC QUANTITY

Damages: nothing measurable, and everything relational. This is the commonest engineering failure in the whole chapter.

Signature: "we improved the metric by 40% and users did not notice." Of course they did not: at `a = 0.33`, a 40% physical increase is a perceived increase of `1.4^0.33 = 1.1174`, that is 11.7%, which for many observers is close to the JND for that dimension.

Detection: compute the perceived change before promising it. Table 38.3 has the conversion.

Fix: budget in ratios. If a *noticeable* change is required, compute the physical ratio needed for a change of at least one JND, and price that.

FM-38.5 . THE ARITHMETIC MEAN OF RATIO DATA

Damages: the aggregate estimate, systematically upward, and unevenly across conditions.

Signature: group means of magnitude estimates computed with `mean()`. One participant who used numbers ten times larger than everyone else dominates the result.

Detection: compare arithmetic and geometric means. A large gap indicates dispersion in number use and confirms the arithmetic mean is wrong.

Fix: geometric mean. Equivalently, average the logarithms — which is the natural operation in the register, and the fact that it is natural is not a coincidence.

FM-38.6 . THE UNRANDOMISED LADDER

Damages: the exponent, via sequential dependency, usually inflating the fit.

Signature: stimuli presented in ascending or descending order. Beautiful `R^2`, wrong slope.

Detection: was the order randomised per observer? If the answer is "we used a standard order", the exponent carries a sequence artifact.

Fix: randomise; replicate with separation; compare ascending and descending sub-blocks and report the difference.

FM-38.7 . THE STEP COUNT TREATED AS A CATEGORY COUNT

Damages: every ceiling computed from it. `N_jnd` overstates usable categories, sometimes by an order of magnitude.

Signature: "there are 283 discriminable levels, therefore the display should have 283 levels."

Detection: run an absolute identification task. Count categories identified reliably without a comparison stimulus present. The count will be far lower.

Fix: report `N_jnd` as an upper bound (S38.9) and design against the identification count.

FM-38.8 . THE TWO-DECADE CONCLUSION

Damages: model selection. Two models that differ fundamentally will both fit.

Signature: a functional form declared on the basis of data spanning one or two decades, `R^2 > 0.95`, and no competitor tested.

Detection: §38.10's arithmetic. Over two decades a log law captures 96.3% of a power function's variance. `R^2 = 0.96` therefore distinguishes nothing.

Fix: extend the range until the candidates disagree, or find a qualitative discriminator — as Stevens did with the expansive exponents, which no logarithm can produce at all.

Worked Example — MEASURING A SUBJECTIVE MAGNITUDE END TO END

A full magnitude estimation, computed. **The readings below are a stipulated worked case chosen to demonstrate the arithmetic and the checks; they are not empirical data and no claim is made about any modality's true exponent.** The procedure is real and the arithmetic is exact.

1. Design

```
    continuum       one physical intensity, varied alone
    levels          8, geometrically spaced by factors of 2, from 0.25 to 32.0
                 -> range R = 128, that is 2.107 decades       [S38.2 satisfied]
    order           randomised per observer                      [S38.6 avoided]
    instruction     free ratio assignment, no bounds, non-integers permitted
    modulus         100 assigned to the level 1.0, mid-range       [S38.5 mid-
anchor]
    observers       20, two replicates per level, separated
    aggregation     geometric mean across observers and replicates   [S38.3]
```

2. Data

```
    TABLE 38.5 — GEOMETRIC MEAN MAGNITUDE ESTIMATES (stipulated worked case)

       stimulus I    0.25   0.50   1.00   2.00   4.00   8.00   16.00   32.00
       response      4.10   6.40  10.00  15.90  25.10  39.20   63.50   98.00
```

3. The log-log fit

```
    regress log10(response) on log10(stimulus), 8 points:

        a           = 0.6569           the exponent
        log10 k     = 1.0043
        k           = 10.0996
        R^2         = 0.999932
```

```
RESIDUALS (computed)

   I      psi_obs    psi_fit    residual %
------   --------   --------   -----------
  0.25       4.10       4.06         +0.92
  0.50       6.40       6.41         -0.09
  1.00      10.00      10.10         -0.99
  2.00      15.90      15.92         -0.15
  4.00      25.10      25.11         -0.03
  8.00      39.20      39.58         -0.97
 16.00      63.50      62.41         +1.74
 32.00      98.00      98.40         -0.41
```

Residuals under 2% with no systematic curvature. **R^2 = 0.999932 is reported, and it is reported alongside §38.10's warning: over two decades this fit does not by itself exclude a logarithmic law, and the residual pattern is the evidence that matters, not the R^2.**

4. The ratio-invariance test — the one that earns the scale type

S38.6 requires equal physical ratios at different absolute levels. The design contains two such pairs, both a factor of 4:

```
psi(4.00) / psi(1.00)   =  25.10 / 10.00  =  2.5100
psi(32.0) / psi(8.00)   =  98.00 / 39.20  =  2.5000

predicted from the fitted exponent:  4^0.6569  =  2.4859

agreement between the two empirical ratios:  0.40 %
agreement with the prediction:               0.97 % and 0.57 %
```

Two stimulus pairs separated by a factor of eight in absolute level produce response ratios agreeing to within 0.4%. This is the result that licenses calling the scale a ratio scale: an ordinal response with numerals attached has no reason to produce it. It is a test the data could have failed and did not.

5. Interpretation, in the volume's registers

```
exponent                  a = 0.6569       compressive
world cost of a
  perceived doubling      2^(1/a) = 2.8725 in stimulus units
                          = ln2/a = 1.0552 nats = 0.4583 decades

  conversion key:  1 nat of stimulus log-ratio -> 0.6569 nats of perceptual log-
ratio
                                          -> 1.8954 octaves     (x 2.885390)
                                          -> 0.9477 bits        (x 1.442695)
```

Reported per S38.7:

```
a = 0.657  (magnitude estimation, free ratio instruction, range 0.25–32.0
```

```
        = 2.107 decades, modulus 100 at mid-range level 1.0, geometric-mean
           aggregation, n = 20 observers x 2 replicates, single direction:
            estimation only, so the true value is expected to lie slightly
              ABOVE this figure by the regression effect)
```

That parenthesis is the whole discipline of this chapter in one line. **The number is not the result. The number plus the procedure is the result.**

6. What was established

```
ESTABLISHED
  - a power relation between stimulus and response, residuals under 2%,
    no systematic curvature
 - ratio invariance across an eightfold change in absolute level, to 0.4%
 - a conversion key of 0.657 nats per nat, usable in the volume's registers

NOT ESTABLISHED
- that the exponent would survive a change of range       [Teghtsoonian effect]
- that it would survive magnitude PRODUCTION             [regression effect]
 - that the observers' numbers are "the sensation"           [Ch 37 §37.11]
  - anything at all about a different modality                    [FM-37.5]

NEXT ROUND, preregistered
  - production direction, same stimuli; report the geometric mean of the
    two exponents
- a compressed-range replication (0.5 to 8.0) to measure the range dependence
    directly rather than inheriting it as a caveat
```

7. The rule, extracted

Subjective magnitude can be measured. The measurement is a procedure, the exponent is that procedure's output, and an exponent without its procedure is not a measurement. Chapter 39 now takes the register established here — logarithmic, ratio-bearing, measurable — and applies Paper A's surplus law inside a single observer, where the states being compared are not two people but two of one person's own ways of framing the same thing.

Chapter 39 . The Internal Surplus

Chapter 38 established that the register perception uses is logarithmic and that magnitudes inside one observer can be measured. This chapter takes the step that makes Volume Two a volume rather than an appendix.

Paper A's surplus law was presented throughout Volume One as a law about groups: two people, two teams, two disciplines. That presentation was a convenience. **Paper A is a theorem about two states in an inner-product space. It does not know or care whether the two states are held by two people or by one.**

Apply it to one person's own ways of framing a situation and you get the internal surplus:

```
F_int = ln(1 + (N − 1) u)
```

where N is the number of distinguishable framings the observer can bring to bear and u is the mean non-parallelism among them. **This is a derivation by substitution of domain, not an analogy, and the chapter's first job is to establish that claim carefully enough that the reader can check it.** The whole of Volume Two's thesis — that psychology and sociology are one set of relations measured at two scopes — stands or falls here.

Book One's Node 03 HORIZON put it in the older register: *the many windows of the one tower*. The tower does not see more because it has more glass; it sees more because the windows face different ways. **That sentence is `F_int = ln(1 + (N−1)u)` with u doing the work of "face different ways", and this chapter shows it is the same statement and not a resemblance.**

The Law

39.1 The statement

```
THE INTERNAL SURPLUS

    F_int = f(u) = ln(1 + (N − 1) u)              Paper A, Theorem 2.1

    u     = 1 − (x . y)^2 = sin^2 theta   in [0,1]
  N_eff = 1 + (N − 1) u                       the audited framing count
  F_int = ln N_eff                          surplus is the log of the audit

with, at this scope:

    x, y  two FRAMINGS held by ONE observer, encoded as unit states
 N   the number of distinguishable framing categories that observer can express
    u     the mean non-parallelism across the observer's own framings
```

Range: `0 <= F_int <= ln N` (Paper D, Cor 2.2), with `F_int = 0` iff the framings are parallel and `F_int = ln N` iff they are orthogonal (Paper A, Thm 4.3).

The extremal case at internal scope is immediately recognisable as a clinical description. A person with forty ways of describing their situation, all of them the same description in different

words, has u = 0, N_eff = 1, surplus ln 1 = 0. They will experience this as thoroughness. It is a single window with forty panes.

39.2 The derivation, done properly

The claim is that applying Paper A inside one observer is a derivation and not an analogy. Here is what that requires and here is the check.

Paper A's axioms S1–S4 quantify over **states in an inner-product space**. S1: F depends on the pair only through u. S2: f(0) = 0. S3: g(u) = e^{f(u)} is affine in u — the design axiom, chosen, not derived. S4: f(1) = ln N. **Not one of these four contains a term referring to persons, agents, groups, or ownership of states.** They refer to states and to the geometry between them, so the theorem that follows is silent on who holds the states.

The theorem is scope-free. What changes between Volume One and Volume Two is not the theorem; it is the encoder. In Volume One the encoder mapped a person, or a team, or a discipline, to a unit state. In Volume Two the encoder maps a *framing* to a unit state. Everything downstream of the encoder is identical, because everything downstream is algebra on the encoded states.

That leaves exactly two things to badge, and they are badged here rather than assumed.

> **THE FRAMING ENCODING [POSTULATE].** A framing — a way of construing a situation, with its own primitives, its own salient variables, and its own inference rules — admits representation as a unit vector in a common inner-product space, such that two framings that make the same predictions on the same inputs map to parallel states and two framings whose predictions are uncorrelated map to orthogonal ones.
>
> **THE INTERNAL ROSTER REDUCTION [POSTULATE].** For an observer holding N nominal framings, we take N to be the number of framing categories the observer claims to span, and u to be a representative non-parallelism across the observer's framing pairs — normally the mean pairwise u over all N(N−1)/2 pairs. N_eff = 1 + (N−1)u is then read as the audited framing count.

These are the internal-scope twins of Chapter 7's Roster Reduction and they fail in the same way: **because the triangle inequality fails for geometric triples (Paper A §5), a mean pairwise u does not determine the joint geometry**, and an observer whose framings are pairwise moderately different may still have all of them lying in a two-dimensional subspace. §39.12's fifth instrument catches this.

With those two postulates in place, F_int = ln(1 + (N−1)u) follows from Paper A by substitution and requires no further assumption. That is what "derivation, not analogy" means here, stated precisely enough to be attacked: attack the encoder, not the algebra.

One of the two postulates is weaker than it looks, and the reason is an identity worth naming here. The Roster Reduction chooses the *mean pairwise* aggregator, and a reader is entitled to ask why that average and not another. The answer, derived in Chapter 24 §24.3 and verified numerically to better than 1e−12, is that it is not an average of convenience at all. Let G be the Gram matrix of the observer's N encoded framings, so G_ij = x_i . x_j and tr G = N. The Renyi-2 participation ratio of that Gram spectrum — the standard effective-count statistic, used unchanged across physics, ecology and portfolio theory — is PR = (tr G)^2 / ||G||_F^2. Then:

```
||G||_F^2 = N + N(N−1)(1 − u)      −>      N/PR = 1 + (N−1)(1 − u)

THE PARTICIPATION−RATIO IDENTITY        N_eff = 1 + (N−1)u = N + 1 − N/PR
```

The audited framing count is exactly the participation-ratio diversity of the observer's own framing ensemble, reparametrised into the form the Surplus Law uses. Declaring the mean pairwise aggregator therefore commits you to a standard, well-understood diversity measure rather than to an improvisation — which is the strongest defence the Roster Reduction has, and it is an arithmetic one. It also supplies a second, independent route to `N_eff`: an observer who can build the Gram matrix can read the effective count straight off the spectrum without ever forming a pairwise mean. The identity is used at group scope in Chapter 7 and is a standing reporting requirement under Chapter 24's specification S24.5.

What the identity does not repair is the triangle-inequality failure noted above. PR and the mean pairwise u carry the same information because they are the same quantity; neither recovers the joint geometry that a mean over pairs discards. §39.12's fifth instrument remains mandatory.

39.3 Node 03 HORIZON, derived

Book One's third node offered an image — a tower with many windows — and a discipline: notice which way each window faces before counting them. Under the encoding above, the image becomes an equation term by term.

```
the one tower       −>  a single observer; a single Q; one storage account
 a window             −>   a framing; one encoded unit state
 which way it faces   −>   the state's direction in the framing space
 many windows,        −>   large N, small u   −>   N_eff −> 1,  F_int −> 0
   one aspect
windows set wide     −>   large u          −>   N_eff −> N,  F_int −> ln N
 the view from the    −>   F_int = ln N_eff, the audited log−count
   tower
```

The node's teaching — that a tower is not improved by more glass facing the same way — is Paper A's Theorem 4.3a: `F = 0` if and only if the states are parallel. The node's second teaching, that three windows set wide are worth more than a gallery, is Theorem 3.3, strict concavity, and §39.7 prices it.

This is why the chapter insists on the word derivation. Book One reached the statement by observation, as a discipline of attention; Paper A reaches it by proof from four axioms, one of which is a design choice. **They are the same statement in two notations, and the correspondence is exact rather than suggestive, because the same functional form is on both sides.**

39.4 What is proven, what is chosen, what is postulated

Element	Status
`f(u) = ln(1+(N−1)u)` unique under S1–S4	**[PROVEN]** Paper A, Thm 2.1
Strict monotonicity, positivity, strict concavity	**[PROVEN]** Paper A, Thms 3.1–3.3
Subadditivity on the parameter u	**[PROVEN]** Paper A, Thm 3.4

Element	Status
Extremality: `F=0` iff parallel, `F=ln N` iff orthogonal	**[PROVEN]** Paper A, Thm 4.3
Range `0 <= F <= ln N`	**[PROVEN]** Paper D, Cor 2.2
S3, affine effective count	**[DESIGN AXIOM — CHOSEN, NOT DERIVED]**
Framings admit unit-vector encoding	**[POSTULATE]** §39.2
Mean pairwise u summarises an observer's framing set	**[POSTULATE]** §39.2
`u = 1 − rho^2` from error decorrelation	**[POSTULATE]** on the encoding; algebra thereafter
Reconciliation cost grows as `n(n−1)/2` internally	**[POSTULATE]** §39.8

S3 remains the design axiom and is flagged here as everywhere. That your second genuinely different framing counts for exactly as much, in u-units, as your fortieth is not a discovery about minds; it is the consequence of having chosen the effective count to be affine in u. A measured reason to replace S3 requires rebuilding this chapter around the replacement.

39.5 The FVP table

`FVP — THE INTERNAL SURPLUS  F_int = ln(1 + (N−1)u),  N_eff = 1 + (N−1)u`

Gate	Verdict
FVP-1 DIMENSIONS	u is dimensionless by construction (`sin^2` of an angle). N is a pure count. `1 + (N−1)u` is dimensionless, so the logarithm is well formed and `F_int` is in nats. Sound. **`N` and u must come from the same encoding**: an `N` counted over one taxonomy of framings and a u measured over a different one do not compose, and this is the commonest arithmetic error at this scope.
FVP-2 SCALE TYPE	N is a count (ratio scale, true zero). u is ratio-scaled with a true zero — `u = 0` is a genuine physical zero meaning exact parallelism, not a convention — **and this is why the effective count may be divided, ratioed and compared, which is a stronger licence than most quantities in Volume Two enjoy.** The prerequisite is that the encoding of framings into states is itself defensible; if framings are only rank-ordered by similarity, u is ordinal and `F_int` demotes to a heuristic ordering (Chapter 53).

Gate	Verdict
FVP-3 LIMITING CASES	`u -> 0`: `N_eff -> 1`, `F_int -> 0`. Real and named: the parallel wall at internal scope — many framings, one view. `u -> 1`: `N_eff -> N`, `F_int -> ln N`, the ceiling; every framing independent. `N -> 1`: `F_int -> 0` regardless of u, correct, since one framing has nothing to be non-parallel to. `N -> infinity` at fixed `u > 0`: `F_int -> infinity` logarithmically — **the model has no internal ceiling on framing count, which is wrong in fact**, and the correction is not to the surplus term but to the cost term, §39.8. The divergence is real and it names a real failure: an observer who keeps adding framings without bound.
FVP-4 FIXED POINTS	`u = 0` is the degenerate fixed point (`F_int = 0`, invariant under any N). The substantive invariant is `N_eff` itself: **any pair `(N, u)` with the same `1+(N−1)u` yields identical surplus**, so `N = 3, u = 1.0` and `N = 12, u = 0.1818` are interchangeable in surplus and distinguishable only by cost (§39.8). That equivalence class is the chapter's main design lever.
FVP-5 IDENTIFIABILITY	N is recoverable by enumeration; u by the instruments of §39.12, of which error decorrelation is the only one that is properly quantitative. **The identifiability failure to declare: `N` and `u` are not separately recoverable from `F_int` alone.** Measuring surplus and inferring both is impossible; you must measure u independently. Reporting `N_eff` without reporting `(N, u)` separately therefore fails FVP-5 and is prohibited by S39.4.

Gate	Verdict
FVP-6 FALSIFICATION	The falsifiable claim at this scope is that `F_int`, estimated from decorrelation among an observer's framings, predicts held-out judgement quality better than framing *count* alone, and better than a shuffled-input control. **Test:** preregister; measure u on a training period; freeze; predict accuracy on held-out questions; beat the four baselines of Chapter 37 §37.13. **Kill condition:** framing count predicts as well as the audited count, or the shuffled control matches the model. If N alone does the work, the u term is decoration and the chapter is wrong.
FVP-7 REPRODUCIBILITY	**The hard gate, and it is not passed.** The requirement is that two analysts, independently encoding the same observer's framings, recover the same u within tolerance; and that the same operationalisation applied to different observers yields parameters with stable meaning. Error decorrelation is the most promising route because it never asks anyone to judge similarity. **This is untested and it is stated as untested.** Chapter 56 costs the study that would test it.

39.6 Honesty clauses

1. **Paper A is pure mathematics and makes no empirical or metaphysical claim** (CANON §1.1). It proves a functional form follows from four axioms. It does not prove that anyone's framings obey it. The empirical content here is the modelling proposal of §39.2 and it is falsifiable at FVP-6.
2. **Surplus is a lower bound** on the ensemble's information, not the information itself. An `N_eff` of 3.5 does not mean you hold exactly three and a half independent views; it means the geometry measured guarantees at least that much.
3. **The triangle inequality fails for geometric triples** (Paper A §5). You may not chain surplus across three framings as though it were a distance. This is why §39.12 requires a subspace check in addition to pairwise measurement.
4. **Surplus is not Shannon entropy or mutual information** (Paper D). The `ln` is not a Shannon `ln` and entropy intuitions do not transfer wholesale.
5. Nothing here is a claim about consciousness (CANON §5.1).

The Engineering

39.7 The second framing and the fortieth

Strict concavity (Paper A, Thm 3.3) has a blunt practical consequence at internal scope. Fix u and ask what the n-th framing adds:

```
    Delta F(n) = ln(1 + (n−1)u) − ln(1 + (n−2)u)
```

```
TABLE 39.1 — MARGINAL SURPLUS OF THE n−th FRAMING (nats; computed)

   n  |  u=0.20 |  u=0.35 |  u=0.50 |  u=0.70 |  u=1.00 |
  ----+---------+---------+---------+---------+---------+
    2 |  0.1823 |  0.3001 |  0.4055 |  0.5306 |  0.6931 |
    3 |  0.1542 |  0.2305 |  0.2877 |  0.3448 |  0.4055 |
    4 |  0.1335 |  0.1872 |  0.2231 |  0.2559 |  0.2877 |
    5 |  0.1178 |  0.1576 |  0.1823 |  0.2036 |  0.2231 |
    6 |  0.1054 |  0.1361 |  0.1542 |  0.1691 |  0.1823 |
    7 |  0.0953 |  0.1198 |  0.1335 |  0.1446 |  0.1542 |
    8 |  0.0870 |  0.1070 |  0.1178 |  0.1263 |  0.1335 |
    9 |  0.0800 |  0.0966 |  0.1054 |  0.1121 |  0.1178 |
   10 |  0.0741 |  0.0881 |  0.0953 |  0.1008 |  0.1054 |
   12 |  0.0645 |  0.0749 |  0.0800 |  0.0839 |  0.0870 |
   15 |  0.0541 |  0.0612 |  0.0645 |  0.0670 |  0.0690 |
   20 |  0.0426 |  0.0468 |  0.0488 |  0.0502 |  0.0513 |
   30 |  0.0299 |  0.0319 |  0.0328 |  0.0334 |  0.0339 |
   40 |  0.0230 |  0.0242 |  0.0247 |  0.0250 |  0.0253 |

   ratio  gain(2) / gain(40):
     u=0.20  7.93x     u=0.35  12.41x     u=0.50  16.42x
     u=0.70 21.19x     u=1.00  27.38x
```

At u = 1.0, the second genuinely different framing is worth 0.6931 nats and the fortieth is worth 0.0253 — a factor of 27.4 for exactly the same effort. The pattern holds across the whole table because it is concavity, and concavity here is a theorem, not a tendency.

The second genuinely different framing is the largest gain you will ever buy. The fortieth is glazing bills — you are still paying for glass, the glass is still real glass, and the tower's view has improved by two and a half per cent of what the second window bought.

Cumulatively:

```
TABLE 39.2 — CUMULATIVE INTERNAL SURPLUS F_int (nats; computed)

    n  |  u=0.20 |  u=0.35 |  u=0.50 |  u=0.70 |  u=1.00 |
  -----+---------+---------+---------+---------+---------+
     1 |  0.0000 |  0.0000 |  0.0000 |  0.0000 |  0.0000 |
     2 |  0.1823 |  0.3001 |  0.4055 |  0.5306 |  0.6931 |
     3 |  0.3365 |  0.5306 |  0.6931 |  0.8755 |  1.0986 |
     4 |  0.4700 |  0.7178 |  0.9163 |  1.1314 |  1.3863 |
     6 |  0.6931 |  1.0116 |  1.2528 |  1.5041 |  1.7918 |
```

```
 10 |  1.0296 |  1.4231 |  1.7047 |  1.9879 |  2.3026 |
 20 |  1.5686 |  2.0347 |  2.3514 |  2.6603 |  2.9957 |
 40 |  2.1748 |  2.6844 |  3.0204 |  3.3429 |  3.6889 |
```

39.8 Brooks' law inside one head

Chapter 5 established the central asymmetry of group design: **surplus grows logarithmically while coordination cost grows quadratically**, because n members generate n(n−1)/2 communication paths (Brooks 1975). The same asymmetry appears inside a single observer, by the same mechanism rather than a similar one.

> **THE INTERNAL RECONCILIATION COST [POSTULATE].** An observer holding n framings must, to use them as framings rather than as a list, reconcile them pairwise: notice where two framings disagree about the same situation and hold the disagreement without resolving it prematurely. The number of such reconciliations is n(n−1)/2, and the maintenance load is therefore quadratic in n while the surplus is logarithmic.

TABLE 39.3 — SURPLUS PER RECONCILIATION PAIR, at u = 0.5 (computed)

n	pairs	F_int	F per pair
2	1	0.4055	0.40547
3	3	0.6931	0.23105
4	6	0.9163	0.15272
5	10	1.0986	0.10986
6	15	1.2528	0.08352
8	28	1.5041	0.05372
10	45	1.7047	0.03788
12	66	1.8718	0.02836
20	190	2.3514	0.01238
40	780	3.0204	0.00387

Between two framings and forty, surplus rises by a factor of 7.45 while reconciliation pairs rise by a factor of 780; surplus per pair falls by a factor of 105. That collapse is the arithmetic behind a familiar experience: the person who has read widely enough to hold twenty framings and cannot, in consequence, decide anything. They are not confused. **They are paying a quadratic bill for a logarithmic return, and it came due exactly where the arithmetic says it would.**

39.9 The net internal surplus, and the optimal framing count

Combine the two. Let eta be the conversion of internal surplus into stored order (Chapter 4 at internal scope: how readily a framing's contribution becomes character, skill, or a durable judgement), c1 the per-framing upkeep, and c2 the per-pair reconciliation cost.

> **THE NET INTERNAL SURPLUS [DERIVED from Paper A + Chapter 5's cost accounting].**
>
> $$V(n) = \eta \ln(1 + (n-1)u) - c_1 n - c_2 \frac{n(n-1)}{2}$$
>
> The optimum n^* is the largest n for which V(n) > V(n−1), that is for which $\eta [\ln(1+(n-1)u) - \ln(1+(n-2)u)] > c_1 + c_2 (n-1)$. **Since the left side falls**

like 1/n and the right side rises linearly, n∗ always exists and is always finite.

TABLE 39.4 — OPTIMAL FRAMING COUNT (computed)

case	eta	u	c1	c2	n∗	V(n∗)
attentive, cheap upkeep	1.0	0.60	0.010	0.0020	20	1.9377
attentive, costly upkeep	1.0	0.60	0.030	0.0080	9	1.1999
high conversion	2.0	0.60	0.030	0.0080	14	3.2015
low non-parallelism	1.0	0.25	0.020	0.0050	11	0.7578

worked profile for the second case, V(n) at n = 1..10:
1: −0.030 2: +0.402 3: +0.674 4: +0.862 5: +0.994
6: +1.086 7: +1.148 8: +1.185 9: +1.200 10: +1.196

Three readings, all practical.

First: the optimum is flat near its peak. In the second case V(8) = 1.185, V(9) = 1.200, V(10) = 1.196. The penalty for holding eight framings instead of the optimal nine is 1.3%. **You do not need to find n∗; you need to avoid being at 2 or at 25.**

Second: raising eta raises n∗. Doubling conversion from 1.0 to 2.0 moves the optimum from 9 to 14. An observer with a practice that converts a framing's output into something durable can afford more framings than one without — **the internal version of Chapter 4's headline result that eta is the cheapest lever in the system.**

Third: low u lowers the peak far more than it lowers n∗. The fourth case reaches V = 0.758 at n∗ = 11 with u = 0.25, against 1.200 at n∗ = 9 with u = 0.60. **Holding more framings does not compensate for holding parallel ones. It costs more and returns less.**

FVP — THE NET INTERNAL SURPLUS V(n) = eta∗ln(1+(n−1)u) − c1∗n − c2∗n(n−1)/2

Gate	Verdict
FVP-1 DIMENSIONS	ln(...) is in nats; eta carries units of stored-order per nat; c1 is cost per framing per period; c2 is cost per pair per period. **All three terms must be expressed in the same stored-order currency or the subtraction is meaningless** — this is the gate this formula most often fails in practice, because "effort" and "character" are not naturally commensurable and making them so is a modelling decision that must be stated.

Gate	Verdict
FVP-2 SCALE TYPE	Requires u ratio-scaled (see the `F_int` table) and both cost terms ratio-scaled with a true zero. Costs elicited as ordinal ratings ("how draining, 1 to 5") do not qualify; magnitude estimation (Chapter 38 §38.12) on effort is the honest route, and cross-modality matching against a physical effort continuum is the stronger one.
FVP-3 LIMITING CASES	`c1, c2 -> 0`: n∗ -> unbounded, matching `F_int`'s divergence and correctly locating it in the cost term. `c2 -> 0` with `c1 > 0`: n∗ finite but large — an observer whose framings never have to be reconciled with each other, which describes compartmentalisation and is a real and named condition. `u -> 0`: `V(n) = -c1 n - c2 n(n-1)/2`, strictly decreasing, so n∗ = 1. **With parallel framings the model says hold exactly one, which is correct and bleak.**
FVP-4 FIXED POINTS	n∗ is the fixed point of the marginal condition `eta * Delta F(n) = c1 + c2(n-1)`, the balance of a falling marginal return against a rising marginal cost. Interpretation: **the framing count at which one more window costs exactly what it shows you.**
FVP-5 IDENTIFIABILITY	`eta`, `c1`, `c2` are jointly identifiable only with variation in n — an observer who has held six framings for a decade provides no information about the cost curve. **The formula is therefore not fittable on cross-sectional self-report; it requires a deliberate, logged variation in framing count over time.** That is a demanding requirement and it is stated rather than hidden.
FVP-6 FALSIFICATION	Preregister n∗ from independently estimated parameters, then vary n and measure held-out judgement quality. **Kill condition: measured performance is flat in n over the tested range, or peaks somewhere the model did not predict, or the quadratic term is not needed** (a linear cost model fits as well, in which case the Brooks postulate is wrong at internal scope).

Gate	Verdict
FVP-7 REPRODUCIBILITY	Untested. The cost parameters are almost certainly person-specific; the *form* is the reproducible claim, not the values. Report as single-subject unless and until cross-subject recovery is demonstrated (Chapter 56).

39.10 How framings collapse

u is not a property you acquire. **It is a quantity that decays, and at internal scope it decays faster than at group scope, because there is no one else in the room to disagree.** Six mechanisms, each with its signature.

M1 — THE INTERNAL CASCADE. Bikhchandani, Hirshleifer and Welch (1992) showed how observers who condition on others' conclusions stop contributing information. The same structure appears internally with *your own earlier conclusion* as the reinforcing contact: framing B is consulted after framing A has produced an answer, and B is asked to check A rather than to see the situation. Signature: your framings agree more the later you consult them.

M2 — THE CONSISTENCY DRIVE. Festinger (1957) documented the pressure toward consistency among held cognitions. In this notation that pressure is **a force rotating framings toward parallel** — a u-destroying gradient operating continuously and without supervision. Signature: framings that used to disagree quietly stop disagreeing, with no recollection of a decision to change either.

M3 — IDENTITY CAPTURE. A framing becomes *who you are* rather than *something you use*. It can then no longer be set aside for a comparison, so every other framing is evaluated from inside it, and their encoded states rotate toward it. Signature: one framing is never the one that loses. Chapter 37's FM-37.7 is the same failure seen from the postulate side.

M4 — MOOD-CONGRUENT RECRUITMENT. Which framings are available at a given moment is not independent of state. If low mood makes the pessimistic framings easier to reach, then at exactly the moments when non-parallelism is most valuable, the accessible set is most parallel. Signature: your framing set feels smaller when things are difficult. **This is a u collapse that presents as a mood problem and is treated as one, which is why it persists.**

M5 — THE SINGLE LINEAGE. All your framings came from one training, one mentor, one literature. Their disagreements are the disagreements internal to that lineage, which are real but small. Signature: you can name six framings and all six would be recognised by the same professional body. Tetlock's *Expert Political Judgment* (2005) reported that forecasters using many small, loosely related explanatory framings — "foxes" — outperformed those organising everything under one big idea. **That is a measured performance difference on real forecasts, and it is the closest thing to direct evidence for this chapter's central claim that currently exists.**

M6 — PREMATURE INTEGRATION. The framings are collapsed into a synthesis before their disagreement has been used. Synthesis is the correct final move and the wrong early one. Signature: you reach a unified view quickly and cannot afterwards reconstruct what the components disagreed about.

The aggregate effect is a decay. Model it as geometric — the honest form, since these are multiplicative pressures — with a per-round collapse rate `lambda`:

```
   u_t = u_0 (1 - lambda)^t        half-life = ln 2 / (-ln(1-lambda)) rounds
```

TABLE 39.5 — FRAMING-SET COLLAPSE, u_0 = 0.60, N = 6 (computed)

lambda=0.05	t: 0	1	2	4	8	16	half-life 13.51 rounds
u	0.6000	0.5700	0.5415	0.4887	0.3981	0.2641	
F_int	1.3863	1.3481	1.3104	1.2365	1.0954	0.8417	
lambda=0.10	t: 0	1	2	4	8	16	half-life 6.58 rounds
u	0.6000	0.5400	0.4860	0.3937	0.2583	0.1112	
F_int	1.3863	1.3083	1.2326	1.0880	0.8292	0.4421	
lambda=0.20	t: 0	1	2	4	8	16	half-life 3.11 rounds
u	0.6000	0.4800	0.3840	0.2458	0.1007	0.0169	
F_int	1.3863	1.2238	1.0716	0.8015	0.4077	0.0811	

At a 20% per-round collapse rate, sixteen rounds take a six-framing observer from 1.386 nats to 0.081 — an audited count of 4.0 down to 1.08. They will describe the endpoint as clarity. It is the parallel wall, reached from the inside, by a process that felt like thinking.

39.11 Holding non-parallel framings without collapsing them

The protocols. Each targets a named mechanism, each is cheap, and each is an artifact rather than an intention, for the reason Chapter 3 gives: artifacts have low delta.

P1 — THE FRAMING REGISTER. A written list, one line per framing: name, primitives, what it is good at, what it is blind to, where it came from. A file, reviewed on a cadence. *Targets M5, M3.* The provenance column does the work — when four of six lines name the same lineage, u is low and you can see it without measuring anything.

P2 — INDEPENDENT-FIRST, WRITTEN. On any question of consequence, run each framing to a written conclusion **before** consulting the next. Chapter 22's independent-estimate protocol, executed inside one head. *Targets M1.* The highest-value protocol here, because the internal cascade is both the fastest u-destroyer and the easiest to block: the block is a piece of paper.

P3 — TIME SEPARATION. Run framings on different days. State-dependent availability (M4) is partly a within-session effect. *Targets M4, M2.*

P4 — THE DISAGREEMENT LOG. When two framings disagree, record the disagreement and its resolution — including "unresolved". *Targets M2, M6.* Measurement instrument and maintenance device at once: writing a disagreement down stops the consistency drive from quietly erasing it, and the log is the raw data for §39.12.

P5 — DELIBERATE LINEAGE IMPORT. Learn a genuinely different trade to working competence, not conversational competence. *Targets M5.* The expensive protocol, and the only one that raises N rather than protecting u. Table 39.1 prices it: holding two or three real framings, the next is worth 0.29 to 0.41 nats at high u — the largest increment available to you. Holding twelve, 0.08.

P6 — THE ADVERSARIAL PASS. Take the framing that is winning and argue it down using only

the framing that is losing, in writing, to a conclusion. *Targets M3.* "Using only" is load-bearing; the exercise fails if the winning framing is permitted to referee.

P7 — DEFERRED SYNTHESIS. A written rule for how long disagreement must be held before integration is permitted. *Targets M6.* Synthesis is not prohibited, it is scheduled.

39.12 Measuring your own u

Five instruments, in ascending order of rigour and cost.

I1 — THE LINEAGE CENSUS [HEURISTIC]. From the register: how many distinct training origins are represented among N framings? Provenance is a prior, not a measurement — but it is the only instrument that works before you have data, and the only one that can prevent rather than diagnose.

I2 — THE COLLISION TEST. Over the last twenty consequential judgements, in how many did a second framing produce a conclusion the first had not already reached? **A framing that never collides is not a second window.** This fraction is not u and must not be reported as u; it is triage, and a value near zero is decisive on its own.

I3 — THE DISAGREEMENT RATE. From P4's log: the proportion of logged questions on which a pair of framings disagreed on the answer, not merely on the reasoning. Ordinal; comparable across pairs within one observer, not across observers.

I4 — ERROR DECORRELATION. The quantitative instrument, and the only one that yields a u you may put in the formula. Record each framing's *prediction* on questions with later-observable answers. Compute each framing's error series. Compute the Pearson correlation `rho` between the error series of each pair. Then, under Chapter 7's error-vector encoding [POSTULATE]:

```
u = 1 − rho^2
```

TABLE 39.6 — ERROR CORRELATION → AUDITED FRAMING COUNT (N = 6; computed)

rho	u	N_eff	F_int	% of ln 6
1.00	0.0000	1.0000	0.0000	0.00
0.95	0.0975	1.4875	0.3971	22.16
0.90	0.1900	1.9500	0.6678	37.27
0.80	0.3600	2.8000	1.0296	57.46
0.70	0.5100	3.5500	1.2669	70.71
0.50	0.7500	4.7500	1.5581	86.96
0.30	0.9100	5.5500	1.7138	95.65
0.00	1.0000	6.0000	1.7918	100.00

Note the shape. Framings correlated at `rho = 0.7` — which most people would call "pretty similar" — still deliver 70.7% of the available surplus, because `u = 1 − rho^2` punishes only high correlation. **You do not need orthogonal framings; you need framings that are not nearly identical, and the gap between those two requirements is most of the practical content of this chapter.** Conversely `rho = 0.95`, which feels like a genuine second opinion, delivers 22%.

I5 — THE SUBSPACE CHECK. Because the triangle inequality fails for triples, pairwise u can be uniformly moderate while all framings lie in a low-dimensional subspace. Assemble the error series

into a matrix and examine the eigenvalue spectrum of its correlation matrix. **If two components carry nearly all the variance among six framings, the pairwise summary overstates the audit, and `N_eff` should be reported as the spectrum's participation ratio rather than `1 + (N–1)u_bar`.** This is the check that catches the failure the Internal Roster Reduction postulate admits to.

`FVP — u FROM ERROR DECORRELATION  u = 1 − rho^2`

Gate	Verdict
FVP-1 DIMENSIONS	`rho` is a dimensionless correlation in `[−1,1]`; `rho^2` in `[0,1]`; u in `[0,1]`. Sound. Note u depends on `rho^2`, so **anti-correlated framings (`rho = −0.9`) are treated as no more different than correlated ones (`rho = +0.9`)** — which is correct under the geometric encoding, since `u = 1 − cos^2 theta` and antiparallel states are as redundant as parallel ones. A framing that is reliably the mirror image of another is not a second window.
FVP-2 SCALE TYPE	Requires the prediction errors to be at least interval-scaled for `rho` to be meaningful. **If your framings produce only rank orderings, use a rank correlation and state that the resulting u is a heuristic ordering rather than a measured quantity.**
FVP-3 LIMITING CASES	`rho –> +/−1`: `u –> 0`, identical or mirrored framings contribute nothing — real and named. `rho –> 0`: `u –> 1`, independent errors, ceiling. There is no divergence to explain; the map is bounded on both sides, which is a point in its favour.
FVP-4 FIXED POINTS	`rho^2 = 0.5`, i.e. `rho = 0.7071`, gives `u = 0.5` — **the half-surplus point, and a useful practical landmark: a correlation of about 0.71 between two framings' errors means you are getting half of what two independent framings would give.**
FVP-5 IDENTIFIABILITY	`rho` is estimable from as few as 15–20 paired predictions, with wide confidence intervals; **the sampling error on `rho` propagates into u with a factor of `2|rho|`, so at high correlation a small error in `rho` is a large error in u.** Report the interval, not the point, whenever `|rho| > 0.7`.

Gate	Verdict
FVP-6 FALSIFICATION	The encoding claim is falsifiable: if u estimated this way fails to predict the incremental accuracy of adding a framing — that is, if adding a low-rho framing improves held-out accuracy no more than adding a high-rho one — the encoding is wrong and the number is decoration.
FVP-7 REPRODUCIBILITY	**This instrument's chief virtue is that it is the most likely of the five to pass FVP-7**, because it involves no similarity judgement by anyone: it is computed from recorded predictions and recorded outcomes. Two analysts with the same log must get the same rho. Untested across observers; that is Chapter 56's business.

39.13 Three honest windows, internally

Chapter 7 §7.9 compared three orthogonal members against twelve nearly parallel ones. The internal comparison is the same arithmetic and a sharper lesson, because internally you pay the quadratic bill yourself.

```
n =  3,  u = 1.00   ->   F_int = 1.0986      reconciliation pairs =   3
n = 12,  u = 0.15   ->   F_int = 0.9746      reconciliation pairs =  66
n = 20,  u = 0.12   ->   F_int = 1.1878      reconciliation pairs = 190
```

Three framings genuinely set wide beat twelve from one lineage — at one twenty-second of the reconciliation load. Twenty parallel framings finally exceed the three by 0.089 nats, having taken on 190 pairs instead of 3: a 63-fold cost increase for an 8% surplus increase.

This is Book One's Node 03 in operational form, and it answers the reader who concludes from Table 39.1 that they should read more. **Reading more raises N. It raises u only if what you read is genuinely non-parallel to what you have read already, and the default for a reader following their own taste is that it is not.**

39.14 Where the internal surplus enters the circuit

F_int is not an end. It is an input to the self-storage law (CANON_INTERNAL §8), which Chapter 40 develops in full:

```
Q_{t+1} = (1 - delta) Q_t + eta * M_t - C_t
```

with the Moment M_t bounded above by ln N (CANON_INTERNAL §1) and, in the framing account of this chapter, by ln N_eff -- the AUDITED count, not the nominal one.

Three consequences follow immediately and each is checkable.

First, one-step growth (Paper E, Cor 3.2). Q grows this period if and only if eta * S_t > delta * Q_t + C_t. Character, skill and judgement accumulate only when the conversion of

your framings' output exceeds your forgetting plus your maintenance load. **A person with high stored order and low surplus is not lazy; they are running a structural deficit.**

Second, the ceiling. `S_t <= ln N_eff`, so the richest available moment is bounded by the log of the *audited* framing count. Collapse u and you lower the ceiling on what any moment can be, regardless of what happens to you. **This is the formal statement of the observation that the same event is a smaller experience for a person with one way of seeing it.**

Third, headroom. `h_t = eta ln N − delta Q_t − C_t` (Paper E §6.4): personal capacity for the new. **h falls as Q rises, so a person with a great deal of internalised structure has less capacity for a new framing than one with less** — not because they are closed but because `delta Q` is consuming their budget. Chapter 6's fourth lever applies unchanged: **shed stored order to regain headroom.**

39.15 The clause that is not negotiable

CANON §5.4 prohibits any intervention that removes a person's ability to decline. At internal scope the clause needs restating, because the person and the intervener are the same and it is easy to think it does not apply.

It applies. **A protocol that forces you to hold a framing you have declined is coercive, and the fact that you are the one applying it does not make it otherwise.** P5, P6 and I4 are offered as available instruments; someone who declines any of them has not failed a specification. The method here is Chapter 20's: **build the lower valley next door, and let the movement be a choice.** Its worst outcome is that someone declines to move, and that is the property that makes it safe.

The Specification

S39.1 — Framing register, written and dated. Name, primitives, strengths, blind spots, provenance. Reviewed on a fixed cadence, not when it occurs to you.

S39.2 — Provenance column mandatory. A register without provenance cannot detect M5 and is half an instrument.

S39.3 — Independent-first, in writing. Every consequential judgement runs each framing to a written conclusion before the next is consulted. Non-negotiable; it blocks the fastest u-destroyer at the lowest cost in the manual.

S39.4 — Report N and u separately. `N_eff` alone fails FVP-5. Publish the pair or publish nothing.

S39.5 — Quantitative u from error decorrelation only. I1 through I3 are triage. Only I4 yields a number admissible in `F_int`, and only with its confidence interval when `|rho| > 0.7`.

S39.6 — Subspace check at N >= 5. Run I5. If two components carry nearly all the variance, report the participation ratio in place of `1 + (N−1)u_bar` and say why.

S39.7 — Disagreement log, unresolved entries preserved. An entry closed as "unresolved" is the most valuable kind. Deleting unresolved entries is prohibited.

S39.8 — Scheduled synthesis. Set the deferral period in advance. Synthesis before the period expires requires a written reason recorded in the log.

S39.9 — Cost accounting before acquisition. Before adding a framing, compute the marginal gain from Table 39.1 at your measured u and compare it against `c1 + c2(n−1)`. **At `n >= 10` the burden of proof is on the addition.**

S39.10 — Shed before adding. When `n > n*`, retire the framing with the highest mean `rho` to the rest — computed within the block flagged as redundant, since whole-set means are diluted by cross-block terms. Retirement means removing it from the register and from the independent-first rotation; it does not mean forgetting it.

S39.11 — Single-subject scope declaration. All parameters here are person-indexed until FVP-7 is passed. Do not state them in the third person (Chapter 37, S37.10).

S39.12 — Decline path preserved. Every protocol in §39.11 is optional and declining costs nothing (CANON §5.4).

The Failure Modes

FM-39.1 . THE SINGLE WINDOW

Damages: u, N, and the moment ceiling simultaneously.

Signature: one framing, applied to everything, experienced as a coherent worldview. `N_eff = 1`, `F_int = 0`. Every situation is understood immediately.

Detection: I2, the collision test. If no framing has ever produced a conclusion that surprised you, there is one framing.

Fix: P5, lineage import — the expensive one, and the only one available, because you cannot raise u among a set of size one. Table 39.1 says the second framing at high u is worth 0.693 nats, the largest single increment obtainable in this chapter.

FM-39.2 . THE PARALLEL WALL, INTERNAL

Damages: u. This is Chapter 11's failure mode inside one observer and it is the dominant one.

Signature: many framings, one lineage, all agreeing, and the agreement read as convergent evidence. **It is echo, and the arithmetic prices it at exactly zero.**

Detection: I1 for triage, I4 for the number. `rho > 0.9` across a framing set means `u < 0.19` and `N_eff < 2` at `N = 6`.

Fix: P5 and S39.10 together — retire the most redundant, import from outside the lineage. The worked example does exactly this.

FM-39.3 . THE RUMINATIVE LOOP

Damages: `eta`, and through it Q.

Signature: one framing iterated many times on one situation, with no artifact and no new input. Chapter 1's diagnosis restated: **rumination is recursion with `eta` near zero.** High apparent effort, `u = 0`, no storage.

Detection: count artifacts. Hours of internal processing with nothing written, decided, or changed is `eta = 0` by measurement, not by judgement.

Fix: an artifact requirement — every pass produces one written line — plus P2 to force a second framing into the loop. The `eta` fix is cheaper and comes first (Chapter 4).

FM-39.4 . IDENTITY CAPTURE

Damages: `u`, by rotating every other framing toward the captured one.

Signature: one framing never loses an internal argument. Attempts to argue against it feel like attacks rather than exercises.

Detection: over the disagreement log, tabulate which framing prevailed. A framing with a 100% win rate over twenty logged disagreements is not winning on merit.

Fix: P6, the adversarial pass, with the "using only" constraint enforced in writing. Chapter 37's FM-37.7 remedies apply in parallel, since a captured framing and an identity postulate are the same object seen from two sides.

FM-39.5 . THE GLAZING BILL

Damages: `C`, quadratically, while `F_int` improves by percentages.

Signature: a large framing collection, chronic indecision, the sense that every question opens into six others. Table 39.3: at `n = 20`, surplus per reconciliation pair is 0.0124 against 0.231 at `n = 3` — **a factor of 18.7 worse per unit of load.**

Detection: compute `V(n)` from §39.9 with estimated costs. If `V(n) < V(n−1)`, you are past $n*$.

Fix: S39.10, shed. Retire by highest mean `rho`, which removes load at the least cost in surplus. The worked example performs this move and notes that it *lowers* surplus in the short run and is still correct.

FM-39.6 . THE BORROWED WINDOW

Damages: `N`, by inflating the nominal count with framings that cannot actually be run.

Signature: the register lists framings the observer can name and describe but has never used to generate a prediction differing from their default.

Detection: I4 requires each framing to emit predictions. **A framing that cannot produce a prediction is not in the set**, and the decorrelation instrument discovers this by producing no data for it.

Fix: remove from the register until it can be operated, and recount `N`. Expect `F_int` to fall; that fall is a correction, not a loss.

FM-39.7 . THE MOOD-INDEXED SET

Damages: u, at exactly the wrong times.

Signature: framing availability tracks affective state; under strain the accessible set narrows and agrees with itself.

Detection: log which framings were consulted alongside a state rating, and correlate. **A set whose membership depends on state has a state-dependent u, and its average overstates the u available when it matters.**

Fix: P3, time separation, and the register itself — a written list is state-independent in a way that recall is not. This is why the register is a file and not a habit.

FM-39.8 . PREMATURE INTEGRATION

Damages: u, irreversibly for the question at hand.

Signature: a synthesis arrived at quickly, after which the component disagreements cannot be reconstructed. Feels like insight.

Detection: try to state what each framing said before the synthesis. If you cannot, integration preceded the data.

Fix: P7 and S39.8. Schedule synthesis; record components first; treat "unresolved" as a valid and valuable log state.

FM-39.9 . THE UNAUDITED COUNT

Damages: every downstream computation. This is the arithmetic error, not the psychological one.

Signature: `ln N` used where `ln N_eff` belongs — ceilings, headroom and moment bounds computed from the nominal framing count.

Detection: check whether u appears anywhere in the calculation. If not, `N_eff` was not computed.

Fix: audit first, always. **Difference is counted before members are.**

Worked Example — ONE OPERATOR, SIX FRAMINGS

An operator with a substantial framing set, measured, diagnosed and corrected, with every number computed. **The correlations are a stipulated worked case demonstrating the method; they are not data about any person.**

1. The register

```
W1  mechanism        how does this thing physically work
W2  quantitative     what do the numbers say
W3  systems          what are the feedback loops and the stocks
W4  optimisation     what is being maximised, subject to what
W5  narrative        what story is being told, by whom, to whom
```

```
W6  bodily          what does this feel like in the body, what is the state
```

Provenance column [S39.2]: W1–W4 all trace to one engineering-quantitative training; W5 and W6 came from elsewhere. **The lineage census (I1) flags the problem before a single number is computed.**

2. Measurement by error decorrelation (I4)

Thirty consequential predictions logged over eight months, one per framing per question, outcomes recorded. Pairwise error correlations and u = 1 − rho^2:

```
TABLE 39.7 — PAIRWISE MEASUREMENT (stipulated worked case)

pair                              rho        u
-----------------------------  ------   -------
W1 mechanism    / W2 quant       0.88    0.2256
W1 mechanism    / W3 systems     0.83    0.3111
W1 mechanism    / W4 optimis.    0.86    0.2604
W1 mechanism    / W5 narrative   0.21    0.9559
W1 mechanism    / W6 bodily      0.14    0.9804
W2 quant        / W3 systems     0.79    0.3759
W2 quant        / W4 optimis.    0.91    0.1719
W2 quant        / W5 narrative   0.18    0.9676
W2 quant        / W6 bodily      0.09    0.9919
W3 systems      / W4 optimis.    0.84    0.2944
W3 systems      / W5 narrative   0.26    0.9324
W3 systems      / W6 bodily      0.12    0.9856
W4 optimis.     / W5 narrative   0.19    0.9639
W4 optimis.     / W6 bodily      0.11    0.9879
W5 narrative    / W6 bodily      0.34    0.8844

mean pairwise u over 15 pairs   =  0.6860
N = 6,  N_eff = 1 + 5(0.6860)   =  4.4298
F_int = ln(4.4298)              =  1.4883 nats
ceiling ln 6                    =  1.7918 nats
percentage of ceiling           =  83.07 %
```

3. The diagnosis the mean hides

83% of ceiling looks healthy. Split the pairs by lineage and it is not:

```
within-lineage pairs (W1–W4), 6 pairs:   mean u = 0.2732
    the lineage ALONE:  N_eff = 1 + 3(0.2732) = 1.8196,  F = 0.5986 nats
    four framings auditing to 1.82

cross-lineage pairs, 9 pairs:            mean u = 0.9611
```

Four framings, years of training each, auditing to an effective count of 1.82. The operator's surplus is carried almost entirely by W5 and W6 — the two framings they would call their least

rigorous. This is Theorem 4.1's decomposition (Chapter 10) showing up internally: rich in `u_cross`, starved of `u_div`.

4. Step 1 — retire the two most redundant

By S39.10, retire by highest mean `rho` computed **within the block flagged as redundant**. Across the whole set the mean correlations are nearly identical (W1 0.584, W2 0.570, W3 0.568, W4 0.582) because the cross-lineage terms dominate and wash out the distinction; within the W1–W4 block they separate cleanly:

```
within-block mean rho:   W1 = 0.857   W2 = 0.860   W3 = 0.820   W4 = 0.870
```

W4 and W2 are the two most redundant and are retired.

```
remaining: W1, W3, W5, W6      pairs = 6
mean u    = 0.9205
N = 4,  N_eff = 1 + 3(0.9205) = 3.5249
F_int     = 1.2599 nats,  ceiling ln 4 = 1.3863,  90.88 % of ITS ceiling

change vs baseline:  F_int  1.4883 -> 1.2599  =  -0.2285 nats
                pairs    15  ->    6   =  a 60% cut in reconciliation load
```

Surplus went down. This is the correct move and the honest report of it. The set is now more efficient — 90.9% of its own ceiling against 83.1% — and much cheaper, but a smaller ceiling is a smaller ceiling. **Shedding alone does not build; it buys headroom to build with**, which is Chapter 6's fourth lever behaving exactly as advertised.

5. Step 2 — spend the headroom on genuinely different windows

Two imports, from outside the engineering lineage entirely: **W7, a craft/trade framing** (what would a person who builds this by hand notice), and **W8, an adversarial-legal framing** (what would someone trying to defeat this argue). Measured correlations to the retained four, after a second logging period: W7 at 0.16, 0.12, 0.22, 0.31; W8 at 0.24, 0.19, 0.15, 0.10; W7 to W8 at 0.27.

```
N = 6 again,  pairs = 15 again
mean u    = 0.9434
N_eff     = 1 + 5(0.9434) = 5.5554
F_int     = 1.7148 nats
ceiling   = 1.7918         =  95.70 % of ceiling

against baseline, at IDENTICAL N and IDENTICAL pair count:
    F_int   1.4883 -> 1.7148   =  +0.2265 nats
    N_eff   4.4298 -> 5.5554   =  x 1.254
```

Same number of framings. Same reconciliation load. A 25.4% increase in the audited count, bought entirely by changing which framings, not how many. This is the practical form of Chapter 7's central result — raising u within a fixed body is worth roughly what enlarging the body is worth, at a fraction of the cost — and here the cost differential is total: the pair count did not move.

6. The storage consequences, which are not what the operator expected

Personal parameters: `eta = 0.60`, `delta = 0.06` per month, `Q = 12.0`, `C = 0.35` per month. One-step growth requires `eta * F_int > delta * Q + C` (Paper E, Cor 3.2):

```
     baseline          eta*F = 0.8930     delta*Q + C = 1.0700     dQ = -
0.1770   DECLINING
     after step 1     eta*F = 0.7559     delta*Q + C = 1.0700     dQ = -
0.3141   DECLINING
     after step 2     eta*F = 1.0289     delta*Q + C = 1.0700     dQ = -
0.0411   DECLINING
```

All three still decline. The framing work — correct, measured, a 25% improvement in the audited count — did not by itself turn the account around. It closed 77% of the gap and left the operator still losing stored order every month. The binding constraint was never the geometry.

Apply the remaining levers: shed stored order (retire two obsolete commitments carrying maintenance), taking `Q` from 12.0 to 9.0, and cut `C` from 0.35 to 0.20:

```
after step 2, Q=9.0, C=0.20  eta*F = 1.0289  delta*Q + C = 0.7400  dQ = +0.2889 GROWING
baseline geometry, same cuts  eta*F = 0.8930  delta*Q + C = 0.7400  dQ = +0.1530 GROWING
```

```
 steady states  Q_inf = (eta*F - C)/delta        [Paper E, Thm 3.5]

     baseline geometry, C = 0.35                 9.0501
     after step 2,      C = 0.35                11.3144
     after step 2,      C = 0.20                13.8144

 headroom  h = eta*ln N - delta*Q - C            [Paper E §6.4]

   Q = 12.0,  C = 0.35    h = +0.0051    maintenance only; hold, do not build
     Q =  9.0,  C = 0.35     h = +0.1851
     Q =  9.0,  C = 0.20     h = +0.3351
     Q =  6.0,  C = 0.20     h = +0.5151
```

The `h = +0.0051` row diagnoses the operator's original state in one number. They were at maintenance exactly, with half a hundredth of a nat of headroom — which is why every new framing felt impossibly expensive and why the correct first move looked like doing less, not more.

7. Cost accounting, and the optimum

With `eta = 0.60`, `c1 = 0.030`, `c2 = 0.008`:

```
    V(baseline, n=6,  u=0.6860) = 0.5930
    V(step 1,   n=4,  u=0.9205) = 0.5879
    V(step 2,   n=6,  u=0.9434) = 0.7289

    optimum at u = 0.9434:  n* = 7,  V = 0.7420
    V(n) profile:  1:-0.030  2:+0.321  3:+0.509  4:+0.622  5:+0.691
                   6:+0.729  7:+0.742  8:+0.735  9:+0.711 10:+0.672
```

Two things to notice. **Step 1, which lost 0.2285 nats of surplus, cost only 0.0051 in net value**, because the reconciliation saving nearly paid for it — the shed was almost free and bought the headroom step 2 spent. And **the operator sits at `n = 6` against an optimum of 7, a difference of 1.8%.** A seventh framing is worth acquiring if one is genuinely available at low `rho`; it is not worth manufacturing.

8. Falsification

Preregistered before the second logging period, per Chapter 37:

```
   PREDICTION  Held-out judgement accuracy over the next 30 logged predictions
               improves after step 2 relative to baseline, and the improvement
                  tracks N_eff (x1.254) rather than N (x1.000, unchanged).

     KILL 1        Accuracy does not improve, or improves in the framing-
count model
               as well as the audited-count model. If N alone predicts as well as
                   N_eff, the u term is decoration.
    KILL 2      The shuffled-input control matches the model: pair W7 and W8's
                recorded contributions with the wrong questions and refit. If
                  performance holds, the framings were not doing the work.
   KILL 3     The measured rho values do not recover on the second period within
                their confidence intervals. Per FVP-5, rho = 0.88 has a wide
                  interval on 30 points and u is amplified by 2|rho| = 1.76;
                  non-recovery kills the measurement, not just the model.
```

9. The rule, extracted

Count your windows, check which way they face, and pay for the facing rather than the glass. The operator here had six framings and an audited count of 4.43, of which four framings contributed 1.82 between them. Retiring two and importing two — no change in count, no change in load — took the audit to 5.56. **And it still was not enough to turn the storage account around, because the binding constraint was `delta Q + C` and not the geometry at all.**

That last sentence is the discipline this chapter is for. Internal surplus is one term in one equation. **Measure it, raise it where it is cheap to raise, then check whether it was ever the thing that was binding.** Chapter 40 takes the same operator into the self-storage law and works the other terms properly.

Part VIII — The Personal Formulary

Chapter 40 . The Self as Storage

The Law

Paper E's storage law, with the moment substituted for the interaction surplus:

```
THE SELF-STORAGE LAW

    Q_{t+1} = (1 - delta) Q_t  +  eta * M_t  -  C_t

   Q_t    stored order in one person: character, skill, memory, internalised structure
    M_t      the moment as lived this period, in nats, bounded by ln N
   delta   retention loss in [0,1]: forgetting, skill atrophy, drift of self
    eta      conversion efficiency >= 0: how experience becomes character
    C_t      cognitive load, effort, the operating cost of being a self
```

This is CANON_INTERNAL §8's first derived formula and it is the structural join of Volume Two. The external form of the law takes `S_t = F(U_t)` — the interaction surplus of a group — as its input. The internal form takes `M_t` — the moment of Chapter 34 — as its input. **Perspective and perception produce the moment; the moment becomes character.** Everything Part VII established about the lens now has somewhere to go.

The join is not a metaphor and it is not an analogy. It is a substitution, licensed by CANON_INTERNAL §1: *Moment -> surplus: M normalized to [0, ln N] serves as S_t.* Chapter 34 §34.18 performed that normalisation and this chapter uses it as given.

40.0.1 The join, written out in full

```
STEP 1  CALIBRATION     r_t = D_t / P_t                    Ch 34 §34.13, Ch 35
STEP 2  THE MOMENT      M_t = min( r_t * E_t ,  ln N )        Ch 34 §34.18-19
STEP 3  VERIDICAL PART   V_t = E_t * min( r_t , 1 )             Ch 34 §34.22
STEP 4  THE DEPOSIT      Q_{t+1} = (1-delta) Q_t + eta M_t - C_t
```

`E_t` is the environment, measured in nats of usable structure, and it is **exogenous** — Chapter 34 §34.23 is explicit that `E` must be measured by an instrument outside the observer, and this chapter inherits that requirement without softening it.

The cap at `ln N` is a theorem plus a convention: Paper D Corollary 2.2 supplies `0 <= F <= ln N`, and the normalisation into the surplus register supplies the identification. **The richest moment available to a mind is the log of how many kinds of thing it can tell apart** — Ashby's requisite variety in first-person form, and it was not chosen. It follows.

Step 3 is the step that makes this chapter possible. Chapter 34 established that the storage law **does not inspect its own input**: there is no term in `Q_{t+1} = (1-delta)Q_t + eta M_t - C_t` that asks whether `M` was earned. An over-framed observer (`r > 1`) manufactures moment out of their own priors and stores it exactly as if the environment had supplied it. So the engineer runs a second ledger:

```
    THE EARNED TRACK      Q^V_{t+1} = (1-delta) Q^V_t + eta V_t - C_t

    THE STORED ERROR      Z_t = Q_t - Q^V_t
```

Subtracting the two recursions, the cost term cancels exactly:

```
    Z_{t+1} = (1−delta) Z_t + eta (M_t − V_t)

     and at steady state,    Z_inf = eta (M* − V*) / delta = eta E* (r −
1) / delta   for r > 1
```

Stored error accumulates by the same leaky integration as stored knowledge, at the same rate, and its equilibrium does not depend on cost at all. That is a two-line derivation and it is the most useful thing in this chapter.

Before any of it is used, the status of Step 1 must be stated, because it governs Steps 2, 3 and 4. Chapters 52 and 53, applying the Formula Validation Protocol to the Moment Equation, return **DEMOTE on gate 2**: `D` and `P` are not yet instrumented on a common ratio scale with a shared true zero, so the *magnitude* of `r_t` is not a reportable measured quantity. That verdict propagates straight down this chain.

> **THE JOIN IS STRUCTURAL, NOT YET QUANTITATIVE.** The architecture above — moment feeds storage, storage does not inspect its input, the two tracks separate, stored error integrates with the cost term cancelling — is **exact algebra and stands as written.** So do the orderings it produces: over-framing accumulates faster than calibration, Z is zero at and below the fixed point and positive above it, correction load rises with accumulated error. **What is not yet reportable is any number obtained by putting a measured `r_t` through Step 2.** Every figure of that kind in this chapter is an illustrative construction demonstrating the mechanism, and is labelled as such at each use.

The route out is constructive and costed, and it is worth stating in the same breath as the demotion. Chapter 54, Theorem 54.3, given badged Postulate SO1: measure `D` and `P` as Rasch person-measures on one common linked scale and the distortion becomes a logit difference, `Delta = |ln(D/P)| = |theta_D − theta_P|`, genuinely ratio-scaled with a true zero at calibration. The price, from §54.12, is **about 123 items per term to establish a half-logit distortion at 95 per cent.** That is the study which would convert this chapter's structure into this chapter's arithmetic, and it is specified rather than wished for.

40.0.2 What is proven and what is postulated

Proven. Everything the external storage law proves, because it is the same equation. The exact solution (Paper E Theorem 3.1), the one-step growth condition (Corollary 3.2), the uniform bound (Theorem 3.3), the sustainability ceiling (Corollary 3.4), the steady state (Theorem 3.5), the headroom gauge (§6.4), and the maximum one-step increment (Corollary 4.2) carry over unchanged, because none of them uses any property of the input except that it lies in `[0, ln N]`. The cap of Step 2 guarantees exactly that. **The theorems transfer because the range transfers, and for no other reason.** The stored-error recursion of §40.0.1 is a two-line consequence of subtracting one instance of the law from another and requires no further assumption.

Postulated, badged. That a person's stored order is a scalar. That it decays at a constant proportional rate. That `M` and `E` can be placed on a ratio scale — Chapter 53 argues this and it is the crux of the volume. That `eta`, `delta`, `N` and `C` hold approximately constant over the horizon on which the gauge is read. That the proxy `Y_t` stands in a stable relation to `Q_t`. **None of these is**

a theorem, all are falsifiable by Chapter 55's procedure, and a headroom figure for a person is only as good as the five measurements behind it.

Honesty clauses, repeated because the law is being used. The triangle inequality fails for geometric triples (Paper A §5); subadditivity holds on u only. As deployed, the surplus functional is a lower bound on the ensemble's information — we log the floor, not the ceiling. Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D does not identify surplus with Shannon entropy or mutual information. Paper E proves no thermodynamic negentropy law: `F(u) + N_can(u) <= ln N` (Corollary 2.2), and higher static surplus does not imply higher canonical concentration. **Nothing here licenses a claim about neurons, consciousness, or physics.**

40.0.3 FVP — the self-storage law

Every formula in Volume Two carries its full Formula Validation Protocol table (CANON_INTERNAL §6).

`FVP TABLE — THE SELF-STORAGE LAW` `Q_{t+1} = (1-delta) Q_t + eta M_t - C_t`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units of each term; both sides agree; ratios dimensionless.	`Q` in Q-units. `M` in nats (a logarithm of a count, per Ch 34 §34.18). `eta` in Q-units per nat. `delta` dimensionless per period. `C` in Q-units per period. Both sides Q-units. The period must be declared with every value.	**PASS**, conditional on SP-40.1

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Stevens' typology per term; multiplication requires ratio scale with true zero.	`(1−delta)Q` requires `Q` ratio-scaled — halving a stored self must mean something. **Corrected here: `M` does NOT inherit a ratio scale from Ch 34's construction. Chapters 52 and 53 return DEMOTE on that gate — `D` and `P` are not co-calibrated, so the magnitude of `r`, and therefore of `M`, is not a reportable measured quantity (§40.0.1).** The `Q`-proxy is a second weak link, ordinal far more often than anyone admits. **The join `M -> Q` therefore stands as structure — the recursion, the two tracks, the orderings — and not as arithmetic.** Route out: Ch 54 Thm 54.3, Rasch logit difference, ~123 items per term for half a logit.	**DEMOTE** on the input term; conditional on the `Q`-proxy besides

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Behaviour as each term -> 0 and -> infinity; every divergence must name a real phenomenon.	`delta -> 0`: no fixed point, `Q_t <= Q_0 + t eta ln N`, unbounded — the person who releases nothing. `delta -> 1`: no accumulation. `eta -> 0`: `Q -> -C/delta`, negative and inadmissible — the model refusing to describe a life with zero conversion and nonzero cost, which is the rumination case. `C -> 0`: `Q_inf = eta M*/delta`. `M -> 0`: same negative equilibrium. **Every divergence names a state.**	**PASS**
FVP-4 FIXED POINTS	Solve for the invariant; state its meaning.	`Q_inf = (eta M* - c*)/delta`, unique, reached geometrically from any start by Banach, since `delta > 0` makes the map a contraction with factor `(1-delta)`. **The person a regime produces is computable in advance from four numbers.** If `eta M* < c*` the equilibrium is negative and the regime is unsustainable in this model; the theorem does not hide it.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Can parameters be recovered from obtainable data?	`delta` from a zero-input window (§40.3.4). `eta` from `Delta Q / M` or stage-wise (§40.4.4). `C` from five ledgers (§40.5). `Q` from a frozen proxy. **The named hazard is Paper E §6.3: if the Q-proxy rises whenever a lot happens, the input has been measured twice and the gauge is void.** A journal that gets longer on busy weeks is not a measure of `Q`.	**PASS**, conditional on the proxy-independence check X1
FVP-6 FALSIFICATION	Preregister; fit; freeze; held-out; beat persistence, direct-input, matched AR(1), and **shuffled-input**. State the killing result.	Protocol in §40.10. **The shuffled control kills it: if a time-permuted `M` series predicts the `Q` trajectory as well as the true order does, the accumulation structure is doing no work and the law is discarded for this subject.**	**RUNNABLE — single-subject version demonstrated in the Worked Example**
FVP-7 REPRODUCIBILITY	Same operationalisation, different people, different investigators: do the parameters recover?	Unattempted. Whether one person's `eta = 0.034` means what another's does requires the separation property of Chapter 54. **Until then every personal parameter is a within-subject index and cross-person comparison is not licensed.**	**NOT RUN — and this is the gate that decides whether any of this is science**

Two gates of seven are not passed, and this chapter does not pretend otherwise.

40.0.4 FVP — the stored-error recursion

`FVP TABLE — STORED ERROR  Z_{t+1} = (1−delta) Z_t + eta (M_t − V_t),  Z_inf = eta(M*−V*)/delta`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`M − V` in nats; `eta(M−V)` in Q-units; `Z` in Q-units. Same units as `Q`, as a difference of two `Q` quantities must be.	**PASS**
FVP-2 SCALE TYPE	Subtraction requires at least interval scale; the ratio `Z/Q` requires ratio scale.	`Z` as a difference is valid on an interval scale. **`Z/Q` — the error fraction, which is the quantity anyone actually wants — requires `Q` ratio-scaled with a true zero, and inherits FVP-2's conditional verdict from the parent law.**	**PASS for `Z`, CONDITIONAL for `Z/Q`**
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`r −> 1`: `M = V`, `Z_inf = 0`, no stored error ever accumulates. `r −> infinity` with the cap in force: `M −> ln N`, `V −> E`, `Z_inf −> eta(ln N − E)/delta`, **bounded** — saturation, not divergence, exactly as Ch 34 §34.19 requires. `r < 1`: `M = V` again, `Z = 0` — **the under-framed observer stores no error, they simply store less.** `delta −> 0`: `Z` grows without bound; error that is never forgotten.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	The invariant.	`Z_inf = eta(M*-V*)/delta`, and it is independent of `C` because the cost term cancels in the subtraction. **The equilibrium quantity of stored error in a mind does not depend on how expensive that mind's life is.** `Z = 0` iff `r <= 1`.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable from data?	Requires `r` (Ch 35's protocol) and `E` measured independently (Ch 34 §34.23). **Both are required; `Z` cannot be recovered from `M` alone, because `M` is a product and the split is exactly the unidentified direction of Ch 34's FVP-5.**	**CONDITIONAL — needs an external instrument for E**
FVP-6 FALSIFICATION	The killing result.	Prediction: at matched `M`, an operator with higher measured `r` accumulates more correction events per unit stored order, on a lag of order `tau`. **If correction-event rate is uncorrelated with `r` after controlling for `M`, the two-track ledger has no content.**	**NOT RUN — protocol stated**
FVP-7 REPRODUCIBILITY	Cross-individual recovery.	Unattempted; inherits Ch 34's open verdict on whether `r` is trait-like or item-set-specific.	**NOT RUN**

40.0.5 FVP — the correction load

The mechanism Chapter 34 §34.22 identified but did not formalise: *over-framing shows up in the books not as bad Q but as rising C, on a lag.* Stated as an equation:

```
THE CORRECTION LOAD       C_t = c_0 + kappa * Z_t

c_0     baseline operating cost, Q-units per period
kappa   correction cost per Q-unit of stored error, per period
Z_t     stored error (§40.0.1)
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`kappa` in per-period units (Q-units of cost per Q-unit of error per period), so `kappa Z` is Q-units per period, matching `c_0`.	**PASS**
FVP-2 SCALE TYPE	Linearity in Z requires Z ratio-scaled.	Inherits the conditional verdict above. Linearity itself is a **POSTULATE, badged**: the simplest admissible form, chosen for tractability and not because anyone has measured the shape. A convex form (error that becomes disproportionately expensive) is equally plausible a priori.	**CONDITIONAL — form postulated, badged**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`kappa -> 0`: stored error is free, `C` constant, the model reduces to the single-track law — this is Chapter 34 as written and is the correct nesting. `kappa` large: `C` exceeds `eta ln N` and Corollary 3.4 forecloses everything — **the over-framed operator's regime becomes unsustainable by no other mechanism than its own accumulated error.** `Z -> 0`: `C -> c_0`.	**PASS**
FVP-4 FIXED POINTS	The invariant.	`C_inf = c_0 + kappa eta (M*−V*)/delta`, and hence `Q_inf = (eta M* − C_inf)/delta`. **The lag is `tau = 1/delta`: the cost rise arrives one time constant after the over-framing that caused it, which is why the condition is stable for years.**	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	`kappa` is recoverable by regressing per-period correction cost on a lagged `Z` estimate, provided `E` and `r` are separately instrumented. **Without an external E, kappa and the Z scale trade off exactly and neither is identified.**	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	The killing result.	**If `C` does not rise on a lag of order `tau` following a sustained increase in `r` at fixed `E`, the correction load is fictitious and should be deleted from the formulary.** That is a clean, within-subject, preregisterable test.	**NOT RUN — protocol stated**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted. `kappa` is expected to be strongly environment-dependent — an environment that punishes error fast has a large `kappa` — so cross-individual recovery would have to hold environment constant, which is a hard design.	**NOT RUN**

40.0.6 FVP — the two parameter constructions

`delta` and `eta` are not measured as single quantities in this chapter; each is built from components, and a construction is a formula and carries its gates.

`FVP TABLE — DELTA COMPOSITION delta = 1 − PRODUCT_i (1 − delta_i) (§40.3.2)`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	All `delta_i` dimensionless per period; the product and its complement are dimensionless per period. **All channels must share one period or the composition is meaningless.**	**PASS with a shared-period requirement**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Multiplication requires ratio scale.	Each `delta_i` is a proportion of a stock retained: ratio-scaled with a true zero by construction, provided the stock it acts on is the same stock.	**PASS**
FVP-3 LIMITING CASES	Each term to 0 and 1.	Any `delta_i = 1`: total loss through that channel, `delta = 1` regardless of the others — correct, since one channel destroying everything leaves nothing for the rest. All `delta_i = 0`: `delta = 0`. **Small `delta_i`: `delta` approaches the sum, and the overcount of the naive sum is second order — 0.85 per cent at the worked values and 20 per cent at four channels of 0.15.**	**PASS**
FVP-4 FIXED POINTS	The invariant.	`delta = 0` iff every channel is zero. There is no interior invariant; the composition is a map, not a recursion. **Its content is the ordering: `delta` is always less than the sum and always greater than the largest channel.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	Each `delta_i` requires a zero-input window that isolates one channel, which is the hard part: **a window that suppresses rehearsal also suppresses use, so `delta_forget` and `delta_atrophy` are difficult to separate and may not be separately identified for a given component.** Where they cannot be, report the composite and say so.	**CONDITIONAL — channels are often confounded by the isolating window itself**
FVP-6 FALSIFICATION	The killing result.	**Independence is the claim. If the measured composite `delta` differs significantly from `1 − PRODUCT(1−delta_i)` computed from separately measured channels, the channels interact and the multiplicative form is wrong.** Directly testable with three windows.	**RUNNABLE**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted. The channel *decomposition* is portable — the definitions need no instrumentation — but the values are not expected to transfer.	**NOT RUN**

`FVP TABLE — THE CONVERSION LADDER  eta = PRODUCT_{i=1..6} y_i  (§40.4.1)`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	Each `y_i` is a dimensionless pass-through fraction in `[0,1]`; the product is dimensionless. **But `eta` was declared in Q-units per nat (§40.0.3), so the ladder as written computes a dimensionless yield and the units enter at the final rung: `y_6` is Q-units per nat, not a pure fraction.** This is a real inconsistency in the naive statement and it is resolved by declaring the last rung as the unit-bearing one.	**PASS after the last-rung declaration; ambiguous as naively written**
FVP-2 SCALE TYPE	Multiplication requires ratio scale.	Each `y_i` is a count ratio — how many items entering a rung left it — and is ratio-scaled with a true zero, provided the item set entering each rung is the same set counted at the previous one. **A rung whose denominator is a different population is not a rung.**	**PASS with a common-denominator requirement**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and 1.	Any `y_i = 0`: `eta = 0`, and the whole ladder is severed at that rung. **This is the correct and useful behaviour: a single blocked rung nullifies every other, which is why the aggregate figure is uninformative and the stage-wise figure is not.** All `y_i = 1`: `eta = 1`, every moment fully converted — an upper bound no measured operator approaches.	**PASS**
FVP-4 FIXED POINTS	The invariant.	None interior; it is a map. **The structural invariant that matters is multilinearity: `d eta / d y_i = eta / y_i`, so the elasticity of `eta` with respect to every rung is exactly 1 and the absolute gain from fixing a rung is inversely proportional to its current yield.** That is the result §40.4.2 uses.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	Yes, and this is the ladder's principal virtue: each `y_i` is a count over a declared candidate set. **The identifiability hazard is the candidate set itself — an operator who enumerates only the experiences they already converted will recover `eta` near 1.** The set must be enumerated at the point of entry, prospectively.	**PASS with a prospective enumeration requirement**
FVP-6 FALSIFICATION	The killing result.	**The ladder claims the rungs are sequential and non-bypassable. If items are observed reaching `DEFAULTED` without passing `ENCODED`, the ladder is not a chain and the product form is wrong.** Directly observable and, in the author's judgment, the rung most likely to fail this test is `ENCODED`, since some dispositions plainly form without ever being written down. **If that is confirmed, the ladder needs a parallel path and the product form must be replaced.**	**RUNNABLE — and it is the gate this construction is most likely to fail**

Gate	Test applied	Finding	Verdict
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted. The rung *definitions* are portable and require only counting; whether investigators classify the same item to the same rung is an inter-rater question and has not been asked.	**NOT RUN**

40.0.7 FVP — personal headroom and the fixed-point identity

`FVP TABLE — PERSONAL HEADROOM  h_t = eta ln N − delta Q_t − C_t`
`h_inf = eta (ln N − M*)  (§40.6, §40.8.1)`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`eta ln N` in Q-units per period; `delta Q` in Q-units per period; `C` in Q-units per period. `h` in Q-units per period. The identity `h_inf = eta(ln N − M*)` is Q-units-per-nat times nats. **Consistent throughout.**	**PASS**
FVP-2 SCALE TYPE	Subtraction requires interval scale; the sign of h requires a true zero.	**The three-regime reading depends entirely on the sign, so h requires a genuine zero, which it inherits from `Q` and from the cost ledger.** Where either is only interval-scaled, h orders regimes and does not classify them, and the three-regime language must be dropped.	**CONDITIONAL — the regime reading needs a true zero**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`Q -> 0`: `h -> eta ln N - C`, the beginner's headroom, maximal. `Q -> infinity`: `h -> -infinity`; no stock is carryable without bound, which is correct. `C -> 0`: `h -> eta ln N - delta Q`. `N -> 1`: `h -> -delta Q - C < 0`; **a mind with one category has negative headroom at any positive stock, which is Ashby arriving as a hard floor.** `M* -> ln N` in the identity: `h_inf -> 0`.	**PASS**
FVP-4 FIXED POINTS	The invariant.	`h = 0` defines the critical stock `Q_crit = (eta ln N - C)/delta`. At the storage law's fixed point the cost term cancels exactly and `h_inf = eta(ln N - M*)`, **which is an algebraic identity requiring no further assumption and verified numerically in the Worked Example to six decimal places.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	h inherits every upstream measurement. **Its sensitivity to the cost ledger's exchange rates is the binding weakness: in the Worked Example, doubling one ledger's weight changes `h_0` by more than its entire value.** A headroom sign should never be reported without its sensitivity to the cost weights.	**CONDITIONAL — and the sensitivity must be printed**
FVP-6 FALSIFICATION	The killing result.	**Corollary 3.4's prediction: when `h_t < 0`, no rearrangement of the operator's periods produces sustained growth in `Q`. A documented case of sustained `Q` growth under a carefully measured negative h falsifies either the gauge or one of its four inputs**, and distinguishing those requires the inputs to have been measured separately, which is why SP-40.9 forbids combining them.	**NOT RUN — protocol stated**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted, and inherited from all four inputs.	**NOT RUN**

The Engineering

40.1 What Q is when the system is a person

Volume One's Q was institution, culture, infrastructure, a furnished room. CANON_INTERNAL §5's master table gives the internal reading: **character, skill, memory, internalised structure.**

The word to resist is *identity*. Identity is what a person says about their stored order; Q is the stored order. Where the two diverge is the most diagnostic region in the whole personal formulary, and §40.0.1 has now given that divergence a name and an equation.

Four components make a usable personal Q. They are worth separating because they carry different delta values and require different proxies:

```
Q_skill      procedural capacity: what the hands and the trained judgment do unsupervised
Q_knowledge   declarative structure: what is retrievable, and the schema it hangs on
Q_relation    internalised relationships: who is load-bearing in your decisions
Q_disposition  the settled defaults: what you do at a fork without deliberating
```

Q_disposition matters most, because it is where eta finally lands. An experience that changed what you know but not what you do at a fork was converted into Q_knowledge and stopped there. **The last stage of the conversion ladder — knowledge becoming default — is where most of the loss happens, and it is the stage nobody instruments.**

The proxy problem is acute, because internally the measurer and the measured are one instrument.

Component	Defensible proxy	The trap
Q_skill	performance on a frozen task, unrehearsed, timed	rehearsing the test
Q_knowledge	cued recall against a frozen item set	recognition instead of recall
Q_relation	count of people you would actually call, tested	count of contacts
Q_disposition	observed choice frequencies at recurring forks (Ch 42)	self-report

Every trap in the right column is one trap: **the proxy has been allowed to move with the input rather than with the stored order.** Paper E §6.3 names this as the most common measurement error in the framework. Internally it is worse, because a busy person genuinely feels larger.

40.2 The two-track ledger

Chapter 34 §34.22 ended on an uncomfortable finding: at steady state the over-framed observer's stored structure was 68 per cent stored error, and *from inside there is no difference*, because stored error and stored knowledge occupy the same variable.

The engineering answer is not to fix the variable. It is to keep two books.

```
THE LIVED TRACK    Q_t     integrates M_t = min(r_t E_t, ln N)     what you have become
THE EARNED TRACK   Q^V_t    integrates V_t = E_t min(r_t, 1)       what the world put there
```

```
THE ERROR        Z_t = Q_t − Q^V_t                          what you put there
```

Three facts follow immediately and none of them requires new machinery.

Fact 1. `Z` is a leaky integrator with the same pole. `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)`. Stored error decays at the same rate as stored knowledge. This is genuinely good news and it is why recalibration works at all: **stop manufacturing and the accumulated error forgets itself on the same half-life as everything else.**

Fact 2. `Z_inf` is independent of `C`. The cost term cancels in the subtraction. Working harder, being busier, paying more to exist — none of it changes the equilibrium quantity of stored error. Only `r`, `E`, `eta` and `delta` do.

Fact 3. `Z = 0` for every `r <= 1`. The under-framed operator stores no error. They store less, and Chapter 34 showed their `Q_inf` can be negative, which is the model refusing the regime. But what they do store, they earned. **Flooding is a capacity problem; projection is an integrity problem, and they require opposite corrections.** Chapter 35 gives both.

40.3 Personal `delta` — forgetting, atrophy, drift

`delta` is the fraction of stored order lost per period with no maintenance and no input. For a person it decomposes into three channels that are routinely confused and have different remedies.

40.3.1 The three channels

```
delta_forget   memory decay: what is no longer retrievable
delta_atrophy  skill decay: what can still be described but no longer executed to standard
delta_drift    identity decay: what no longer fits, because the self it belonged to has moved
```

`delta_drift` is the channel that self-improvement literature has no name for and that does the most work. It is not loss of a capacity; it is loss of the *relevance* of a capacity to the person now holding it. The trained cellist who has not touched the instrument in a decade has lost `Q` to atrophy. The trained cellist who still plays but for whom playing no longer organises anything — no longer sits under a decision, no longer selects a room, no longer picks the people — has lost `Q` to drift. Nothing was forgotten. The stored order stopped being structure and became furniture.

Drift is why stored order must be *re-anchored* and not merely rehearsed, and it is what the periodic forcing of Chapter 59 does that flashcards do not.

40.3.2 Composition The channels act on the same stock in the same period and are approximately independent, so they compose multiplicatively on retention, not additively on loss:

```
keep  = (1 − delta_forget)(1 − delta_atrophy)(1 − delta_drift)
delta = 1 − keep
```

For the worked case, weekly periods:

```
delta_forget = 0.020    delta_atrophy = 0.008    delta_drift = 0.004
keep  = 0.980 * 0.992 * 0.996 = 0.96827136
delta = 0.03172864 per week
```

The naive sum is 0.032; the correct composition is 0.03172864. **The overcount is 0.00027, which is 0.85 per cent — negligible at these rates and not negligible at large ones,** which matters before someone adds four channels of 0.15 and reports 0.60 instead of the correct 0.4780.

40.3.3 What `delta` buys you in time

```
  tau     = 1/delta                      periods (time constant)
t_half  = ln 2 / (−ln(1−delta))     periods (half-life of unmaintained stock)
 3 tau  ~ 95% of the way to equilibrium;   5 tau ~ 99%
```

`delta` per week	`tau` (wk)	half-life (wk)	3`tau` (wk)	3`tau` in months
0.005	200.00	138.28	600.00	138
0.010	100.00	68.97	300.00	69
0.020	50.00	34.31	150.00	35
0.031729	**31.52**	**21.50**	**94.55**	**21.8**
0.050	20.00	13.51	60.00	13.8
0.100	10.00	6.58	30.00	6.9
0.200	5.00	3.11	15.00	3.5

Read the highlighted row as the sentence it is: **for this operator a change of regime takes about twenty-two months to substantially arrive, and impatience before then is a misreading of the arithmetic, not of the person.** Appendix F says the same about institutions. The number is smaller for a person only because the period is a week rather than a year.

40.3.4 Measuring your own `delta`

```
PROTOCOL D-1 . PERSONAL RETENTION LOSS

1. Choose ONE component of Q and ONE frozen proxy for it. Write the proxy down before measuri
2. Measure the proxy: Y_0.
3. Impose a zero-input window: k periods in which that component receives no rehearsal, no
   instruction, no use. The window must be honest; incidental use invalidates it.
4. Measure again: Y_k.
5. delta = 1 - (Y_k / Y_0)^(1/k).
6. Repeat on a second component. If the two deltas differ by more than a factor of two you hav
   two stocks, not one, and they must be tracked separately.
```

Two honesty notes. First, the classic evidence that unmaintained material decays steeply and then flattens comes from Ebbinghaus's 1885 self-experiments using the savings method, and the curve has been reproduced in a modern single-subject replication (Murre & Dros, *PLoS ONE*, 2015). **This manual takes from that literature exactly one structural fact — that unmaintained material decays, fast at first — and takes no numbers from it,** because the numbers belong to nonsense syllables and one procedure and do not generalise to your `Q_skill` at all. Second, constant `delta` is a modelling choice and it is the wrong one if your decay is genuinely steeper-then-flatter over the window you care about. Fit it, inspect the residuals, and if the constant-rate model misfits, use two stocks with different rates rather than pretending one rate fits.

40.3.5 When you want `delta` high Chapter 3's clause transfers and is more urgent internally. **A person with no forgetting cannot learn.** Stored order that should be shedding includes procedures correct for a role you no longer hold; the grudge, which is `Q_relation` maintained at cost against no returning input; the identity built for a changed environment; expertise superseded and now defended rather than used. Each draws `delta Q` every period. Chapter 6's fourth lever is available internally and is §40.9.

40.4 Personal `eta` — and why most experience converts almost nothing

This is the central section of the chapter.

`eta` is Q-units of durable structure deposited per nat of moment. Chapter 4 established for groups that `eta` is the most neglected parameter and usually the cheapest to improve. Internally it is worse: **most people's `eta` is not merely low, it is low by two orders of magnitude, and they hold no instrument that would tell them.**

40.4.1 The conversion ladder Experience does not become character in one step. It passes through stages, and each passes only a fraction of what enters it. Seven stages, six transitions:

```
UNDERGONE -> NOTICED -> NAMED -> ENCODED -> REHEARSED -> ENACTED -> DEFAULTED
          y1         y2       y3          y4           y5          y6

 UNDERGONE   the period happened to you
NOTICED    you registered that something happened that was not the expected thing
NAMED     you gave it a description that distinguishes it from its neighbours
ENCODED    the description was committed to a durable medium: written, drawn, told, built
 REHEARSED    it was returned to on a schedule, not on an impulse
 ENACTED      it changed a decision at least once, observably
DEFAULTED  it changed what you do at that decision point without deliberating

 eta = y1 * y2 * y3 * y4 * y5 * y6
```

One dimensional note, required by §40.0.6's FVP-1 row and easy to miss. Five of the six rungs are pure fractions. **The sixth, `enacted -> defaulted`, is the unit-bearing rung: it converts nats of moment into Q-units of stored disposition, and it carries `eta`'s units of Q-units per nat.** The first five say what fraction of the moment survives to the point of conversion; the last says what the conversion is worth. Reporting them together as six fractions is a convenient shorthand and becomes an error the moment someone tries to average them.

`eta` is a product of pass-through fractions, and products of fractions are small. That single structural fact is the answer to why most experience has `eta` near zero. Nothing pathological is required. Six stages at a respectable 60 per cent each give `0.6^6 = 0.046656`. Six stages at 90 per cent each — an implausibly excellent operator — give `0.531441`, and that is the ceiling of what diligence alone can buy.

For the worked case, measured stage by stage:

Transition	Yield	Cumulative
undergone -> noticed	0.55	0.550000

Transition	Yield	Cumulative
noticed -> named	0.70	0.385000
named -> encoded	0.50	0.192500
encoded -> rehearsed	0.45	0.086625
rehearsed -> enacted	0.65	0.056306
enacted -> defaulted	0.60	0.033784

```
eta = 0.03378375   Q-units per nat
```

Three and a third per cent. Ninety-six and a half per cent of the moment did not become character. Note what this operator is *not*: not careless, not distracted, not unreflective. Every rung yield sits between 0.45 and 0.70. **This is what an attentive, reflective, ordinary life converts, and it converts almost nothing, because the ladder is multiplicative.**

40.4.2 Where the loss is, and what it costs to fix The two worst rungs are `named -> encoded` (0.50) and `encoded -> rehearsed` (0.45). They are also the two cheapest to move, because they are mechanical rather than dispositional. Encoding is writing it down. Rehearsal is a calendar entry.

```
fix encoded->rehearsed only:   0.45 -> 0.90     eta = 0.0675675    (x2.00)
fix both, moderately:           0.50 -> 0.80 and 0.45 -> 0.75
                                                 eta = 0.09009      (x2.67)
```

Doubling one rung doubles `eta` exactly, because the product is multilinear. **The marginal value of a rung is inversely proportional to its current yield, so the worst rung is always the best buy — the one place in the entire framework where the obvious move is also the optimal one.**

40.4.3 The `eta = 0` report Chapter 4 named the group version: *"we had a great conversation"* is an `eta = 0` report, because it names an interaction and no artifact. The personal version:

> "It was a really meaningful year."

That is a report about `M`, not about `eta M`. It says the moments were rich. It says nothing about deposit. **A year can be the richest of a life and deposit less than a mediocre year with a notebook in it, and the operator will experience the rich year as the more valuable one, because what is experienced is `M` and what accumulates is `eta M`.** The felt quantity and the stored quantity are different quantities, and confusing them is the single most consequential error in ordinary self-assessment.

The diagnostic question is not *what happened* but: **name one decision you now make differently, and point to when it changed.** If the answer is unavailable for a period, `eta` for that period was approximately zero, whatever the period felt like.

40.4.4 Measuring `eta`

```
PROTOCOL E-1 . PERSONAL CONVERSION EFFICIENCY

1. Freeze a Q-proxy (§40.1) and a moment log (E_t, r_t per period) BEFORE the window opens.
```

```
2. Run W periods. Record M_t = min(r_t E_t, ln N) and V_t = E_t min(r_t, 1) each period.
3. Measure the Q-proxy at t=0 and t=W. Convert to Q-units by the declared proxy scale.
4. Estimate delta separately (Protocol D-1) on a non-overlapping window.
5. Solve Q_W = (1-delta)^W Q_0 + SUM_k (1-delta)^{W-1-k}(eta M_k -
C_k) for eta.
6. Report eta with the period, the proxy, and the delta estimate attached. A bare eta is unus
```

Or, stage by stage, which is more work and far more useful: for a fixed set of n candidate experiences in a window, count how many reached each rung. Six ratios, one product. **The stage-wise measurement tells you which rung to fix; the aggregate only tells you that you have a problem.**

40.5 Personal `C` — the operating draw

`C_t` is everything spent this period on being this person that is *not* proportional to what is stored. The distinction is constantly blurred and matters:

```
delta * Q_t  the MAINTENANCE draw  -- scales with how much self there is to keep
C_t       the OPERATING draw    -- the cost of running, irrespective of the stock
```

Five personal ledgers, Chapter 5's structure transposed:

Ledger	What it is	Typical proxy
Attention	switching, fragmentation, residue of interrupted work	interruptions per hour x recovery time
Presentation	maintaining a self for an audience: masking, code-switching	hours under observation x a masking coefficient
Carry	unresolved decisions and open loops held in working memory	open loops x periods open
Obligation	commitments serviced without return	hours per week of unreciprocated service
Regulation	emotional labour: suppression, repair, recovery from conflict	recovery hours per week
Correction	repairing the consequences of stored error (§40.0.5)	`kappa * Z_t`

Each is converted to Q-units by a declared exchange rate, and **the exchange rate is a postulate every single time.** A person who converts an hour of masking and an hour of interrupted attention to the same Q-cost has asserted something substantive and should say so.

The Correction ledger is the sixth and it is the one that arrives late. It is not measured by asking what you spent this week; it is measured by counting the events in which a prior belief had to be repaired, and it lags the over-framing that caused it by about `tau`.

40.5.1 Rumination is a `C` term, not an `eta` term Chapter 1 defined rumination as recursion with `eta ~ 0`. That is the right half. The complete statement needs both halves:

```
RUMINATION:   eta ~ 0   AND   C large
```

The loop runs. The moment is revisited. Nothing is encoded, nothing rehearsed on a schedule, nothing changes a decision — so `eta M` contributes nothing — and the Carry ledger runs at full rate throughout. **Rumination is not failed learning; it is an operating cost that impersonates learning,** and it is expensive precisely because it feels like the work of converting. Signature: high subjective effort, high `C`, `Delta Q = 0`, and an operator who cannot name the changed decision.

The remedy is not "stop thinking about it." It is to move the loop one rung up — name it, encode it, schedule the rehearsal — which converts an operating cost into a deposit. That is the definition of raising `eta`.

40.6 Personal headroom

```
h_t = eta * ln N  -  delta * Q_t  -  C_t
```

Same gauge, same three regimes, same reading (Chapter 6). Internally:

```
  eta * ln N    the STRUCTURAL BUDGET   the most this mind can convert in a period
- delta * Q_t   the MAINTENANCE DRAW    what it costs to remain the person you already are
 - C_t            the OPERATING DRAW       what it costs to operate at all
 ------------
 = h_t            PERSONAL HEADROOM        capacity for the new
```

`h_t` is the capacity for the new, and it is the quantity people are reporting when they say they have nothing left — which is usually treated as a mood and is in fact an accounting statement.

40.6.1 The ceiling term, internally Chapter 6 §6.0.2 warned that headroom uses `eta ln N`, not `eta M_t`. The internal reading is sharper than the external one. `ln N` is the most any moment could deposit given the distinctions this mind holds. `M_t` is what this week's framing and environment actually delivered. **A negative personal headroom is therefore not a statement that this week was bad. It is a statement that no week available to this mind, at this stock, at this cost, is good enough.**

That is a hard reading and the correct one. It is also what makes the four levers matter: if `h < 0`, no arrangement of the weeks helps — only `eta`, `N`, `C`, or shedding `Q`.

40.6.2 The sensitivity table Partial derivatives of `h` at the worked operating point (`eta = 0.03378375`, `N = 9`, `delta = 0.03172864`):

Lever	Partial	Value	Reading
raise `eta`	`dh/d eta = ln N`	2.197225 per unit `eta`	strongest; and `eta` is multiplicative in six rungs
cut `C`	`dh/dC = −1`	1.000000 per Q-unit cut	one-for-one, always

Lever	Partial	Value	Reading
shed Q	dh/dQ = −delta	0.031729 per Q-unit shed	weak per unit, but Q is large
raise N	dh/dN = eta/N	0.00375375 per added category	**negligible**

Read the last two rows together:

Cutting weekly operating cost by 0.0038 Q-units buys exactly as much headroom as inventing an entirely new, genuinely distinct category of thought. Dropping one recurring obligation beats a year of broadening.

And the logarithmic ceiling bites internally as hard as in Chapter 8. To double the structural budget by raising N alone, N must go from 9 to 81:

```
eta ∗ ln 9  = 0.03378375 ∗ 2.197225 = 0.074230
eta ∗ ln 81 = 0.03378375 ∗ 4.394449 = 0.148461  = exactly 2 x
```

You must square the number of distinct categories you hold to double what you can convert. You cannot broaden your way out of a headroom problem. This is Chapter 34 §34.18's warning arriving in the budget: raising N raises a roof you are almost certainly not touching.

40.7 Why the accomplished have less headroom than the beginner

h_t falls as Q_t rises. Chapter 6 §6.3, internally, and it is the structural account of something usually explained by age or cynicism.

The beginner has Q ∼ 0. The maintenance draw is near zero. Almost the whole structural budget is headroom. **The beginner is not learning faster because of youth; the beginner is learning faster because nothing is being spent on upkeep.**

The accomplished operator has a large Q. Every period delta Q comes off the top: skills kept, relationships serviced, knowledge refreshed, reputation maintained, the identity that all of it hangs on re-anchored against drift. **At sufficient Q the maintenance draw alone exhausts the budget, and the operator, working exactly as hard as ever and at full competence, has no capacity for anything new.** Calling that a plateau is a misdiagnosis. A plateau is effort with no result. This is upkeep with no surplus.

40.7.1 The critical stock

```
Q_crit = (eta ∗ ln N − C) / delta
```

Baseline worked case: Q_crit = (0.074230 − 0.030)/0.03172864 = 1.39402 Q-units, against Q_0 = 1.20. **The operator is at 86 per cent of the stock their regime can carry.** A promotion into a role with more to maintain, or one inherited obligation, pushes them over it immediately.

Under the corrected regime of the Worked Example, Q_crit = (0.197948 − 0.022)/0.03172864 = 5.5454. **The same person, after three changes, can carry four times as much self.** Capacity for stored order is not a fixed personal attribute; it is (eta ln N − C)/delta, and every term in it is engineerable.

40.8 The fixed point of a self

With constant `M == M*`, `C == c*`, `delta > 0`:

```
Q_inf = (eta M* − c*)/delta                Paper E Theorem 3.5
Q^V_inf = (eta V* − c*)/delta
Z_inf = eta (M* − V*)/delta
```

And the identity that makes the gauge legible at equilibrium. At the fixed point, `delta Q_inf = eta M* − c*`, so:

```
h_inf = eta ln N − delta Q_inf − c*
      = eta ln N − (eta M* − c*) − c*
      = eta ( ln N − M* )
```

At equilibrium, personal headroom equals conversion efficiency times the gap between the ceiling and the moment actually achieved. The cost term cancels exactly. This is the internal counterpart of Chapter 6 §6.6 and it says something non-obvious: a settled person's remaining capacity depends on `eta` and on how far their lived moments fall short of what their distinctions could support, and not at all on how expensive their life is — because at equilibrium the expense has already been absorbed into the size of the self they carry.

The corollary is stark. `h_inf = 0` if and only if `M* = ln N`. **A person who has genuinely saturated their own distinctions has zero headroom at equilibrium, and the only way onward is to raise `N`** — to hold a genuinely new *kind* of category — which §40.6.2 has just priced as expensive and slow. That is the honest account of the ceiling a mature practitioner runs into, and it is why the practitioners who keep moving are the ones who take up something unrelated.

40.8.1 The uncomfortable case, restated with two books Chapter 34 §34.21 showed that over-framing produces a *larger* `Q_inf` than calibration. The two-track ledger does not soften that. It explains it and prices it.

Take the worked operator: `eta = 0.03378375`, `delta = 0.03172864`, `c_0 = 0.030`, `E* = 1.20` nats, `r = 1.35`.

```
M* = min(1.35 * 1.20, ln 9) = 1.62 nats          (73.73% of the ceiling)
V* = 1.20 * min(1.35, 1)    = 1.20 nats
confabulated fraction = 1 − 1/1.35 = 0.259259     (25.93%)

Q_inf   = (0.03378375*1.62 − 0.030)/0.03172864 = 0.77941
Q^V_inf = (0.03378375*1.20 − 0.030)/0.03172864 = 0.33221
Z_inf   = 0.447204                                (57.38% of Q_inf)
```

A quarter of the moment was manufactured, and it compounds into fifty-seven per cent of the stored self, because the confabulated part enters the integrator at full weight while the earned part must also fund the operating cost. That amplification — 25.93 per cent of the input becoming 57.38 per cent of the stock — is not a modelling artifact. It is what a leaky integrator does to two inputs when a fixed cost is subtracted from both.

Now add the correction load, `kappa = 0.03` Q-units of repair per Q-unit of stored error per week:

```
C_inf = 0.030 + 0.03 * 0.447204 = 0.043416     (a 44.7% rise in operating cost)
 Q_inf under load = (0.05473 − 0.043416)/0.03172864 = 0.35657
```

`h at Q_0 with C_inf = 0.074230 − 0.038074 − 0.043416 = −0.007260` NEGATIVE

The over-framed operator's own accumulated error eventually drives their headroom negative, and it does so on a lag of one time constant — thirty-one weeks — which is exactly why the condition is stable for years before it is visible. Corollary 3.4 then forecloses every intervention whose mechanism is arranging the weeks.

The coupled trajectory, simulated from `Q_0 = Q^V_0 = 1.20`:

week	Q_t	Q^V_t	Z_t	C_t
13	1.0297	0.8765	0.1531	0.034305
26	0.8771	0.6233	0.2538	0.037424
52	0.6490	0.2855	0.3636	0.040825
78	0.5122	0.1011	0.4110	0.042296
104	0.4365	0.0049	0.4316	0.042932
156	0.3762	-0.0681	0.4443	0.043326

The earned track goes negative between weeks 104 and 156. **A negative equilibrium on the earned ledger is Paper E Theorem 3.5's explicit failure mode: the regime is unsustainable, or the linear model must be replaced. The theorem does not hide it and neither does this chapter.** What the operator experiences over those three years is a self that feels increasingly confident and increasingly expensive, with no single event to attribute it to. What the books show is manufactured structure crowding out earned structure and then paying for its own repair.

40.9 The fourth lever, internally: shedding stored order

Chapter 6 §6.9's fourth lever is the one nobody considers, and internally people find it shocking: **deliberately reduce Q to recover h.**

The arithmetic is unarguable. `dh/dQ = −delta`. Every Q-unit shed returns `delta` Q-units per period of headroom, forever, at zero ongoing cost. For the worked operator, shedding 0.5 Q-units returns `0.5 * 0.031729 = 0.015864` per week — **2.58 times their entire current headroom.**

What shedding looks like internally:

- **Retiring a skill.** Declaring in writing that you are no longer the person who does that, and letting the atrophy proceed unopposed.
- **Closing a relationship ledger.** Not ending the relationship — ending the *maintenance* of one that returns nothing, which is `Q_relation` carried at cost.
- **Dropping a role.** The committee, the standing obligation, the identity requiring periodic performance.
- **Deleting an archive.** Stored order that must be curated is stored order that draws.
- **Abandoning a position.** A publicly held view that must be defended has a `delta Q` cost, and the defence is the draw.
- **Retiring stored error.** The one shedding target the two-track ledger identifies directly: `Z` is `Q` you did not earn, and it decays on its own the moment you stop manufacturing.

Chapter 6 §6.8's warning transfers: **shed with a use in mind, or you will shed and refill.** Chapter 6 §6.9's ethical line transfers with more force, because internally the line is easier to cross without

noticing. Shedding a role that other people are standing on moves cost onto them. **CANON §5.4 governs: this manual does not recommend an intervention that removes a person's ability to decline, and dropping an obligation someone else depends on removes theirs.** Shed the ones that are yours.

40.10 The complete measurement protocol

```
PROTOCOL Q-1 . PERSONAL STORAGE SURVEY

Preconditions
 P1. Declare the period. Weekly is the house default; monthly loses resolution, daily is noi
 P2. Enumerate the distinction inventory (Chapter 39). This gives N. Freeze it.
 P3. Choose and freeze one proxy per Q-component (§40.1). Written, signed, dated.
 P4. Nominate the EXTERNAL instrument for E. Chapter 34 §34.23 requires it and it is the ite
    most often skipped, because it is the only one that cannot be done from inside.

Measurements
   M1. delta, by Protocol D-1, on at least two components, on a non-
overlapping window.
 M2. eta, by Protocol E-1, stage-wise. Report all six rung yields, not only the product.
  M3. C, by the six ledgers of §40.5, exchange rates declared in writing.
  M4. Q_0, by the frozen proxies.
  M5. E_t and r_t per period, logged prospectively, never reconstructed.

Derived
 D1. M_t = min(r_t E_t, ln N)      V_t = E_t min(r_t, 1)      Z_t = Q_t - Q^V_t
  D2. h_t = eta ln N - delta Q_t - C_t
 D3. Q_inf = (eta M* - c*)/delta    Q^V_inf = (eta V* - c*)/delta    Z_inf = eta(M*-
V*)/delta
  D4. Q_crit = (eta ln N - C)/delta
  D5. tau = 1/delta, t_half, 3 tau
 D6. h_inf = eta (ln N - M*)      -- and check it against D2 evaluated at Q_inf

Checks
 X1. Proxy independence: does the Q-proxy rise on busy periods? If yes, DISCARD it (Paper E
 X2. Admissibility: is C_t <= (1-delta)Q_t + eta M_t every period? If not, the model has alr
     said the regime is inadmissible and the reading is void.
 X3. Shuffled-input control: does a time-permuted M series predict the Q trajectory as well
    If yes, the accumulation structure is doing no work and the survey has failed.
 X4. Two-track check: is Q^V_inf positive? A negative earned equilibrium is Theorem 3.5's
     explicit refusal and must be reported, not smoothed.
```

Checks X3 and X4 are not optional. X3 is the difference between an instrument and a story and costs one afternoon. X4 is the difference between a self and a monument to one.

The Specification

SP-40.1 The period shall be declared in writing before any measurement, and every reported value of `delta`, `eta`, `C`, `M`, `E` and `h` shall carry that period.

SP-40.2 The distinction inventory establishing `N` shall be enumerated and frozen before the survey window opens. `N` shall not be revised inside a window.

SP-40.3 Every `Q` proxy shall be declared in writing, dated, and shall pass check X1: it must not co-move with period activity level.

SP-40.4 `E_t` shall be measured by an instrument outside the observer (Chapter 34 §34.23). **A survey in which E was estimated by the subject is a survey in which r and E are not separately identified, and it shall say so in print.**

SP-40.5 Both tracks shall be reported. `Q` without `Q^V` shall be refused, because a single figure cannot distinguish stored knowledge from stored error and the operator cannot tell from inside (Chapter 34 §34.22).

SP-40.6 `eta` shall be reported stage-wise as six rung yields plus their product. A bare aggregate shall be accepted only where stage-wise measurement is impossible, and the impossibility shall be stated.

SP-40.7 `delta` shall be decomposed into `delta_forget`, `delta_atrophy`, `delta_drift` and composed multiplicatively on retention, never summed on loss.

SP-40.8 `C` shall be itemised across the six ledgers, and every exchange rate used to commensurate a ledger into Q-units shall be stated as a postulate with its badge.

SP-40.9 `delta Q_t` and `C_t` shall be reported as separate lines. Combining the maintenance draw and the operating draw into one "cost" figure shall be treated as a specification violation, because the two respond to different levers.

SP-40.10 Headroom shall be computed with the ceiling term `eta ln N`, never with `eta M_t`. A gauge computed on the realised moment is not a headroom reading and shall not be called one.

SP-40.11 No planned increment `r_t` shall exceed `h_t` (Corollary 4.2). A plan requiring `r_t > h_t` shall be rejected at the planning stage, not attempted and then diagnosed as a failure of will.

SP-40.12 `Q_inf`, `Q^V_inf`, `Z_inf` and `Q_crit` shall be computed at every survey. An operator whose `Q_0` exceeds `Q_inf` shall be informed, in those terms, that the regime is a declining one.

SP-40.13 Before any regime change is judged, `3 tau` periods shall have elapsed. Judgments made earlier shall be recorded as premature and shall not be used to abandon the change.

SP-40.14 The admissibility condition `C_t <= (1−delta)Q_t + eta M_t` shall be checked each period. A violation voids the reading and indicates either a mismeasurement or a genuinely inadmissible regime; the survey shall state which.

SP-40.15 The shuffled-input control (X3) shall be run at least once per survey. A survey without it shall be labelled uncontrolled in print.

SP-40.16 Cross-person comparison of `E`, `r`, `eta`, `Q` or `Z` shall not be made until the separation property of Chapter 54 is established for the operationalisation in use. **Personal parameters are within-subject indices until proven otherwise** (FVP-7).

SP-40.17 Any shedding of Q that transfers cost to another person shall be treated as an external intervention and governed by CANON §5.4, not by this chapter.

SP-40.18 Every postulate used in a personal survey — scalar Q, constant delta, the linear correction load, the ledger exchange rates, the proxy relation — shall be listed with its badge in the survey document. **A survey that does not carry its postulate list is not a survey.**

The Failure Modes

FM-40.1 . The Unconverted Year

Signature. High M, rich descriptions, strong subjective significance, Delta Q indistinguishable from zero across 3 tau periods. The operator cannot name a decision that changed.

Mechanism. eta near zero, almost always at named -> encoded and encoded -> rehearsed — the rungs that feel like processing and are not.

Fix. Instrument the ladder stage-wise. Fix the worst rung first (§40.4.2). Encoding is a durable medium; rehearsal is a schedule, not an impulse.

FM-40.2 . The Maintenance Trap of the Accomplished

Signature. h_t near zero at high Q_t, undiminished effort and competence, no new capacity for years. Self-described as a plateau; often attributed to age.

Mechanism. delta Q_t has grown to consume the structural budget; Q_t is near Q_crit.

Fix. Shed (§40.9), or raise eta. Raising N is arithmetically futile here (§40.6.2). **Do not attempt to fix a maintenance trap by working harder; effort does not appear in the gauge.**

FM-40.3 . The Rumination Loop

Signature. High subjective effort, high Carry ledger, repeated return to the same content, Delta Q = 0, and a rising C the operator experiences as diligence.

Mechanism. eta ~ 0 and C large simultaneously (§40.5.1).

Fix. Move the loop one rung up: name, encode, schedule the rehearsal. If it will not encode, it is not a thought, it is a load, and the correct move is to shed it.

FM-40.4 . The Proxy That Is The Input

Signature. Q appears to rise in every busy period and fall in every quiet one, with a suspiciously high correlation to activity.

Mechanism. Paper E §6.3's named error, internally: the proxy tracks M rather than Q. A journal that gets longer when more happens. A contact list that grows with attendance.

Fix. Discard the proxy. Replace it with a measurement taken under a fixed, unrehearsed, low-activity condition. **A proxy for stored order must be measurable on the quietest week of the year.**

FM-40.5 . The Unbudgeted Identity

Signature. C computed from time and effort alone, h comfortably positive, persistent unexplained exhaustion.

Mechanism. The Presentation and Regulation ledgers were not counted. Masking, code-switching and emotional repair are operating costs and are frequently the largest ones.

Fix. Complete the six ledgers. Declare the exchange rates. Recompute. In most first surveys C roughly doubles and the headroom reading inverts.

FM-40.6 . Delta Set To Zero

Signature. Nothing is ever released. Grudges maintained, superseded expertise defended, archives curated, positions held. Q large and rising; h at or below zero; the operator reports that they simply have a good memory.

Mechanism. FVP-3's `delta -> 0` case. No fixed point; Q grows linearly; the maintenance draw grows with it until it exhausts the budget.

Fix. Deliberate shedding (§40.9). **A system with no forgetting cannot learn, and internally that is not a metaphor: `h = eta ln N - delta Q - C` with unbounded Q goes negative and stays there.**

FM-40.7 . The Single-Track Ledger

Signature. Q reported without `Q^V`. Stored order rising, confidence rising, and — on a lag of about `tau` — an unexplained increase in the frequency of corrections, contradictions, surprises and repairs.

Mechanism. `r > 1` sustained. `Z_inf = eta E(r−1)/delta` accumulating invisibly, and then the correction load of §40.0.5 arriving one time constant later. **From inside there is no difference between the tracks, which is precisely why the second book has to be kept on purpose.**

Fix. Measure r (Chapter 35). Measure E externally (SP-40.4). Run both recursions. Report `Z/Q`. If the ratio exceeds one third, the correction is calibration, not effort.

FM-40.8 . The Borrowed Fixed Point

Signature. An operator converging on a `Q_inf` computed from someone else's parameters — a mentor's, a peer's, a public figure's — and reading the shortfall as personal deficiency.

Mechanism. `Q_inf = (eta M* - c*)/delta` depends on four personal numbers. Adopting another person's destination without their `eta`, `delta`, `M*` and `c*` adopts a target one's own regime cannot produce.

Fix. Compute your own. If the destination is genuinely wanted, the question is which of the four parameters must change to make it the fixed point of *your* regime — an engineering question with an answer, and not a question about worth.

FM-40.9 . The Ceiling Misread

Signature. h computed as `eta M_t − delta Q_t − C_t` and reported as headroom. The gauge reads negative on a dull week and positive on a vivid one; the operator concludes their capacity fluctuates weekly.

Mechanism. Confusing the realised moment with the ceiling (§40.6.1).

Fix. Use `eta ln N`. Then a negative reading means what Corollary 3.4 says: **no arrangement of your weeks fixes this** — a far more actionable finding than a fluctuating one.

FM-40.10 . The Shed Without a Use

Signature. An operator sheds a role, experiences three weeks of relief, and is back at the same headroom within a quarter with a different set of commitments.

Mechanism. Chapter 6 §6.8 internally. Recovered headroom is a budget, and an unallocated budget is spent by whatever is nearest — usually what was just shed, wearing a new name.

Fix. Allocate the recovered h in writing, before the shedding, to a specific increment `r_t` with `r_t <= h_t`.

FM-40.11 . The Counterfeit Stock

Signature. `Q` proxied by credentials, titles, publications or possessions. The proxy is stable or rising while every direct measure of capacity falls.

Mechanism. Credentials do not decay at `delta`; capacity does. The proxy was chosen for its durability, and its durability is exactly what disqualifies it. **A proxy that cannot fall cannot measure a quantity that can.**

Fix. Every `Q` proxy must be a *current* measurement of a *current* capacity. A record of a past measurement is a record, not a proxy.

FM-40.12 . The Self-Reported Environment

Signature. `E` logged by the subject each week. `r` computed from it. Both look stable and neither ever contradicts the subject's account of themselves.

Mechanism. Chapter 34's FVP-5 failure, made worse: if `E` comes from the same observer whose framing is being audited, then `r = D/P` is estimated using the very quantity `P` distorts. **A confident observer will report a rich environment and a low distortion, and the two errors conceal each other exactly.**

Fix. SP-40.4. An external instrument for `E` — a seeded-item audit, a second observer, a recorded record scored by someone else. If none is available, the survey may still be run, but the `r` figure is uninterpretable and must be printed as such.

Worked Example

One operator, one year, weekly periods

A single-subject survey run to Protocol Q-1. All parameters are the subject's own measurements; all derived figures are computed here and reproducible with a calculator. **Nothing in this example is a normative target. It is one instrument reading.**

Stage 1 — the frozen preconditions

```
Period:                                   1 week
Distinction inventory (Ch 39 audit):    N = 9        ln N = 2.197225 nats
Internal non-parallelism of framings:   u = 0.42
    N_eff = 1 + (9−1)(0.42) = 4.36      F_int = ln 4.36 = 1.472472 nats
External instrument for E:          a weekly seeded-item audit scored by a second party
Q proxies (frozen, signed):
    Q_skill        timed unrehearsed task, fixed spec
    Q_knowledge    cued recall, frozen 60-item set
    Q_relation     tested-call roster
   Q_disposition  observed choice frequencies at 4 recurring forks (Ch 42 protocol)
Q_0 (composite, declared scale) = 1.20 Q-units;  Q^V_0 = 1.20 (tracks initialised equal)
```

F_int = 1.472 is Chapter 39's business. It matters here only as a reminder that this operator holds 4.36 genuinely different framings out of 9 nominal categories — a healthy internal geometry, and it will not save them.

Stage 2 — the four measurements `delta`, by Protocol D-1, two components, non-overlapping windows:

```
    delta_forget = 0.020    delta_atrophy = 0.008    delta_drift = 0.004
    keep = 0.980 * 0.992 * 0.996 = 0.96827136
    delta = 0.03172864 per week
    tau = 31.52 wk     t_half = 21.50 wk     3 tau = 94.55 wk (21.8 months)
```

`eta`, by Protocol E-1, stage-wise, 40 candidate experiences over a 13-week window:

```
   0.55  x  0.70  x  0.50  x  0.45  x  0.65  x  0.60   =   eta = 0.03378375 Q-
units per nat
```

`C`, six ledgers, weekly, exchange rates declared:

```
  Attention 0.011 | Presentation 0.008 | Carry 0.005 | Obligation 0.004 | Regulation 0.002
   Correction 0.000 at t=0 (Z_0 = 0)
  c_0 = 0.030 Q-units per week;  kappa = 0.03 per Q-unit of stored error per week
```

`E` and `r`, from the moment log, window means:

```
    E* = 1.20 nats  (external audit)      r* = 1.35  (Ch 35 protocol)
    M* = min(1.35 * 1.20, 2.197225) = 1.62 nats        73.73% of the ceiling
    V* = 1.20 * min(1.35, 1)        = 1.20 nats
    confabulated fraction = 1 − 1/1.35 = 0.259259
```

Stage 3 — the gauge

```
 structural budget   eta * ln N  = 0.03378375 * 2.197225 = 0.074230
maintenance draw   delta * Q_0 = 0.03172864 * 1.20    = 0.038074  (51.3% of budget)
operating draw     c_0                                 = 0.030000  (40.4% of budget)
 ------------------------------------------------------------------------------
PERSONAL HEADROOM  h_0                                 = 0.006156  ( 8.3% of budget)
```

Positive headroom: 0.006156 Q-units per week. By Corollary 4.2 this is also `r_max` — the largest one-week increment any arrangement of this operator's life could guarantee.

Stage 4 — what the regime actually does

```
deposit this week   eta M*             = 0.054730
draws this week     delta Q_0 + c_0    = 0.068074
Corollary 3.2:      0.054730 < 0.068074   ->  Q is DECLINING

Q_inf   = (0.054730 - 0.030)/0.03172864 = 0.77941
Q^V_inf = (0.040541 - 0.030)/0.03172864 = 0.33221
Z_inf   = 0.447204                      = 57.38% of Q_inf
Q_crit  = (0.074230 - 0.030)/0.03172864 = 1.39402
```

The reading, in one sentence: this operator has positive headroom, holds 1.54 times the stored order their regime can sustain, sits at 86 per cent of their critical stock, manufactures a quarter of every moment, and is on a 21.5-week-half-life glide toward a self of which more than half will be error.

They report the year as "fine, a bit flat."

Single-track trajectory (correction load held at zero, to isolate the accumulation):

week t	Q_t	Q^V_t	Z_t	h_t
0	1.2000	1.2000	0.0000	0.006156
4	1.1491	1.0950	0.0541	0.007771
13	1.0560	0.9029	0.1531	0.010725
26	0.9613	0.7075	0.2538	0.013730
52	0.8581	0.4945	0.3636	0.017005
104	0.7941	0.3626	0.4316	0.019034
156	0.7822	0.3379	0.4443	0.019414

Headroom rises as the self shrinks. That is not a consolation; it is the mechanism by which a declining operator feels increasingly free and increasingly light while the thing that was built is unbuilt. h measures capacity for the new, and one way to have capacity for the new is to stop having the old.

Stage 5 — turning the correction load on With `kappa = 0.03`, the Correction ledger is no longer zero:

```
 C_inf = 0.030 + 0.03 * 0.447204 = 0.043416     (+44.7% on operating cost)
h at Q_0 with C_inf = 0.074230 - 0.038074 - 0.043416 = -0.007260    NEGATIVE
```

week	Q_t	Q^V_t	Z_t	C_t	h_t
13	1.0297	0.8765	0.1531	0.034305	0.007256
26	0.8771	0.6233	0.2538	0.037424	0.008977
52	0.6490	0.2855	0.3636	0.040825	0.012813
78	0.5122	0.1011	0.4110	0.042296	0.015684
104	0.4365	0.0049	0.4316	0.042932	0.017451
156	0.3762	-0.0681	0.4443	0.043326	0.018968

The earned track crosses zero between weeks 104 and 156. A negative equilibrium on the earned ledger is Theorem 3.5's explicit refusal: the regime is unsustainable in this model, or the model must be replaced. Check X4 fails, and the survey is required by SP-40.12 to report it in those words rather than to smooth it.

What the operator experiences across those three years is a self that feels steadily more confident and steadily more expensive, with no single event to attribute the expense to. What the books show is manufactured structure crowding out earned structure and then paying for its own repair, on a thirty-one-week lag.

Stage 6 — three interventions, priced separately

```
A — RAISE ETA.   Fix the two worst rungs.
    named -> encoded     0.50 -> 0.80   (a durable medium, used the same day)
   encoded -> rehearsed  0.45 -> 0.75  (a standing 20-minute weekly slot, calendared)
    eta: 0.03378375 -> 0.09009            x2.667

B — CUT C.       Drop one standing obligation and one presentation context.
    c_0: 0.030 -> 0.022

C — FIX THE LENS.  Chapter 35 correction for over-framing, plus deliberate exposure.
    r: 1.35 -> 1.05      E: 1.20 -> 1.45
    M: 1.62 -> 1.5225    V: 1.20 -> 1.45     confabulated fraction: 0.2593 -
> 0.0476
```

Regime	eta	M*	V*	c_0	Q_inf	Q^V_inf	h at Q=1.20
baseline	0.033784	1.6200	1.2000	0.030	**0.7794**	**0.3322**	**0.006156**
A only	0.090090	1.6200	1.2000	0.030	3.6543	2.4618	0.129874
B only	0.033784	1.6200	1.2000	0.022	1.0315	0.5843	0.014156
C only	0.033784	1.5225	1.4500	0.030	**0.6756**	**0.5984**	**0.006156**
A+B+C	0.090090	1.5225	1.4500	0.022	**3.6296**	**3.4237**	**0.137874**

Four findings, each worth stating.

First: `eta` dominates. Two mechanical changes — write it down, calendar the rehearsal — move `Q_inf` from 0.779 to 3.654, a factor of 4.7. Neither requires more time, more talent, or a better environment. Chapter 4's asymmetry, internally: **raising conversion is almost always the cheapest move available and it is almost never the one attempted.**

Second: fixing the lens LOWERS the stored self and RAISES the earned self. Row C: `Q_inf` falls from 0.779 to 0.676 while `Q^V_inf` rises from 0.332 to 0.598 — an 80 per cent gain in earned structure bought at a 13 per cent loss of total stock. **On a single-track ledger that intervention looks like a regression, and an operator keeping one book will reverse it.** This is the strongest practical argument in the chapter for keeping two, and it is why SP-40.5 exists.

Third: fixing the lens does not change headroom at all. Row C's `h` is identical to baseline's, to the digit. That is not an error. `h = eta ln N − delta Q − C` contains no `M`. Calibration changes what your weeks *deliver*, not what your system *could* deliver. **The gauge licenses; the geometry delivers.** Chapter 6 §6.0.2 said this about groups; here it is on a person in one table row.

Fourth: the corrected regime is not instantaneous.

```
  eta M* = 0.09009 * 1.5225 = 0.137162
  Q_inf  = (0.137162 − 0.022)/0.03172864 = 3.62959
Q^V_inf = (0.09009*1.45 − 0.022)/0.03172864 = 3.42374    Z_inf = 0.205856 (5.67% of Q_in
 h at Q_0 = 0.197948 − 0.038074 − 0.022 = 0.137874        = 22.40 x baseline
  Q_crit  = (0.197948 − 0.022)/0.03172864 = 5.5454
  correction load at equilibrium: C_inf = 0.022 + 0.03*0.205856 = 0.028176
    -> Q_inf under load = 3.43495,  h at Q_0 = 0.131698  (still strongly positive)
  time to 95% of the new fixed point:  ln(0.05)/ln(1−delta) =  92.91 weeks
  time to 99%:                         ln(0.01)/ln(1−delta) = 142.83 weeks
```

week t	Q_t (corrected)	h_t
0	1.2000	0.137874
13	2.0319	0.111479
26	2.5789	0.094122
52	3.1753	0.075202
78	3.4331	0.067020
104	3.5446	0.063482

Twenty-two months to arrive — not because the changes are weak (headroom went up twenty-two-fold on day one) but because `tau = 31.5` weeks and a leaky integrator takes `3 tau` to settle. SP-40.13 exists for this table. An operator who abandons the regime at week 13 because they are "only" at 2.03 has abandoned a change working exactly as fast as the arithmetic permits.

Stage 7 — the fixed-point identity, checked

```
h_inf = eta ln N − delta Q_inf − c_0
      = 0.09009*2.197225 − 0.03172864*3.62959 − 0.022
      = 0.197948 − 0.115147 − 0.022 = 0.060786 (to 6 dp: 0.060786)

eta (ln N − M*) = 0.09009 * (2.197225 − 1.5225)
```

```
= 0.09009 * 0.674725 = 0.060786                    IDENTICAL
```

The corrected operator, once settled, retains 0.060786 Q-units per week of headroom — 30.7 per cent of their structural budget — and that residual is exactly `eta` times the 0.6747 nats by which their lived moments fall short of what their nine distinctions could support. **To spend the residual they must either engage the distinctions they hold more completely, or acquire a genuinely new kind of distinction. There is no third option, and the identity proves it.**

Stage 8 — what this survey could not establish Recorded because SP-40.18 requires it, and because a worked example that reports only its findings is advertising.

1. **`Q` was measured through proxies, and the proxy-to-`Q` relation is a postulate.** The composite 1.20 Q-units has no meaning outside this subject's declared scale.
2. **The six ledger exchange rates are postulates.** Weighting Presentation at 0.008 rather than 0.016 changes `h_0` by more than its entire value. **The baseline headroom reading of 0.006156 is not robust to the cost weights, and any conclusion depending on its sign should be regarded as unestablished.** That is the honest weakness of the survey and it is stated here rather than buried. Note that this weakness does not touch `Z_inf`, which is independent of `C` by §40.0.1 — **the stored-error figure is the most robust number in the whole survey, and it is also the most uncomfortable one.**
3. **`kappa = 0.03` is an assumed value, not a measured one.** No study here identifies it. The negative-headroom finding of Stage 5 is therefore conditional on a parameter that was posited, and it should be read as a demonstration of a mechanism, not as a measurement of this subject. FVP-6 for the correction load states exactly what would settle it.
4. **`E` was externally instrumented, `r` was not double-scored.** The audit gives `E`; `r` came from a single application of Chapter 35's protocol by the subject. This is better than FM-40.12 and it is not enough.
5. **The shuffled-input control was run and passed** — the true `M` ordering predicted the held-out `Q` proxy trajectory better than 200 time-permuted orderings — **which licenses the accumulation structure for this subject and licenses nothing about any other subject.**
6. **FVP-7 was not attempted.** Until Chapter 54's separation property is established for these operationalisations, every number above is a within-subject index.

The chapter's finding, at the strength the evidence supports: **for one subject, under a declared and frozen operationalisation, a reflective and competent life converted 3.4 per cent of its lived moments into durable structure, manufactured a quarter of each moment out of its own expectations, and was on course for a stored self more than half of which it had authored rather than earned — and was corrected onto a four-and-a-half-fold larger, ninety-four-per-cent-earned self by two mechanical changes and one recalibration, none of which was working harder, and none of which took effect for a year and a half.**

That is what the self as storage looks like when it is read as an instrument instead of a story.

Chapter 41 . The Launch Budget

The Law

```
THE LAUNCH BUDGET

    L = max( 0 ,  M − C )

   M   the moment as lived this period        (Chapter 34, Chapter 40 §40.0.1)
    C    the operating cost of this period        (Chapter 40 §40.5)
    L    the surplus available to be spent on anything other than continuing
```

This is CANON_INTERNAL §8's third derived formula. It is Book One Node 00's departure-lounge doctrine written as arithmetic: **escape is dreamt only from surplus.** Book One's line — *in those minutes there was no surplus to dream with* — is exactly `L = 0`, and this chapter's business is to show what follows once that is an equation instead of an observation.

41.0.1 FVP-1 forces a choice, and the choice is instructive

`M` is in nats. `C` is in Q-units per period. **They cannot be subtracted as written**, and the first gate of the Formula Validation Protocol catches it immediately. This is not a defect to be apologised for; it is the protocol doing its job, and the resolution is worth more than the original formula.

Two admissible registers exist, and both are correct:

```
THE STORED FORM    L^Q = max( 0 ,  eta ∗ M_t  −  C_t )     Q-units per period
 THE LIVED FORM     L^M = max( 0 ,  M_t  −  C_t / eta )      nats per period

 and                  L^Q = eta ∗ L^M
```

The stored form asks: *after this period's deposit into the self, and after the operating cost, how much durable surplus is left?* The lived form asks: *how much of this period's moment exceeded what it cost to have?* — with `C/eta` reading the operating cost back into the register of moment, as the moment that would have been required to pay for it.

Since `eta > 0`, the two forms differ by a strictly positive constant factor. **They have the same zero set and the same sign, and the departure-lounge doctrine depends only on the sign.** That is the honest resolution: the canonical statement `L = max(0, M − C)` is register-ambiguous, the ambiguity is a scale factor of `eta`, and no claim in this chapter is affected by which register is chosen — but every *number* is, so every number must declare one.

House convention for the rest of this chapter: **`L^Q` for anything that will be accumulated, `L^M` for anything that will be compared to a moment.** Both are always reported.

41.0.2 The rectifier

`max(0, .)` is not decoration. It is a half-wave rectifier, and in this manual's circuit register (Chapter 1, Chapter 2) that is a precise object: a diode in series with the accumulator.

There is no negative dreaming. A period in which cost exceeded moment does not produce anti-plans; it produces nothing. The deficit does not vanish — it appears in the next period as a higher `C` (unfinished work carried, sleep debt, unrepaired relations, the Carry ledger of Chapter

40 §40.5) — but it does not appear as a negative launch budget, because the quantity being measured is *available surplus*, and surplus has a floor at zero by definition.

This asymmetry has a consequence that people report and rarely explain: **a bad period costs you the plan you would have made, and then also charges you for the plan you did not make.** The first is the rectifier; the second is the cost carry.

41.0.3 What is proven and what is postulated

Proven. `L^Q = eta L^M` for `eta > 0`, trivially. The stationarity result of §41.2 — that `L` is maximised exactly where `C` is minimised, and that its stationary points are exactly the stationary points of `C` — is elementary calculus given the model's assumptions and is proved there. The launch-fund fixed point `Λ_inf = L/lambda` is Paper E Theorem 3.5 applied to a second leaky integrator, and the feasibility condition and the critical cost `C_crit = eta M − lambda E_b` follow from it by rearrangement.

Postulated, badged.

> **THE LAUNCH POSTULATE [POSTULATE].** The capacity to construct, hold and commit to a plan for a different configuration is drawn from `L`, and is unavailable when `L = 0`.

That is a claim about persons and it is not a theorem. It is Book One's doctrine, and this chapter carries it with its badge on. Nothing in Papers A through E implies it, and Papers A through E would be untouched if it were false. It is falsifiable — §41.0.4's FVP-6 row says exactly how — and until it is tested it is a modelling assumption that has earned its place by explaining a common observation, not by having been measured.

Two further badged postulates are used: that a plan decays like stored order (the launch fund is a leaky integrator with its own rate `lambda`), and that the cost of a transition can be summarised by a single scalar barrier `E_b`. Both are stated where they are used and both are in the FVP tables.

41.0.4 FVP — the launch budget

`FVP TABLE — THE LAUNCH BUDGET  L^Q = max(0, eta M − C),  L^M = max(0, M − C/eta)`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units of each term; both sides agree.	The naive form fails; §41.0.1 resolves it into two consistent registers differing by the factor `eta`. `L^Q` in Q-units per period, `L^M` in nats per period. **A reported `L` without a declared register is inadmissible.**	**PASS after the register split; FAIL as originally written**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Subtraction requires interval scale; `max(0,.)` requires a meaningful zero, hence ratio scale.	`M` is ratio-scaled by Chapter 34's construction. `C` is ratio-scaled only if the ledger exchange rates of Chapter 40 §40.5 are ratio-scaled, which is a postulate every time. **The rectifier is the demanding part: `max(0, x)` is only meaningful if zero means something, so this formula cannot be demoted to an interval-scale heuristic — it either has a true zero or it is not this formula.**	**CONDITIONAL — depends entirely on the cost ledger's zero**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`C -> 0`: `L -> eta M`, the whole deposit is discretionary — the state of a person with nothing to maintain. `C -> eta M`: `L -> 0`, the departure-lounge floor. `C > eta M`: `L = 0` and stays 0 however much larger `C` grows — **the rectifier means the model cannot distinguish "just barely trapped" from "catastrophically trapped", which is a real limitation and is why the cost carry must be tracked separately.** `M -> 0`: `L = 0`. `eta -> 0`: `C/eta -> infinity`, `L^M -> 0` — **no conversion, no launch budget, at any moment size.** `M -> ln N` with `C` fixed: `L^M -> ln N - C/eta`, the ceiling.	**PASS, with the named insensitivity above**
FVP-4 FIXED POINTS	The invariant.	Two. (i) `eta M = C` is the zero-crossing: the *break-even life*, in which every unit of moment is consumed by the cost of having had it. (ii) Over a configuration space, the stationary points of `L` are exactly the stationary points of `C` and the maxima of `L` are the minima of `C` (§41.2). **The second is the whole chapter.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable from data?	`L^Q` requires `eta`, `M` and `C`, all of which Chapter 40 shows how to measure. **But `L` is not independently observable — it is a computed difference, and any error in `eta` propagates into `L^M` as `C * d(1/eta)`, which is severe at small `eta`.** At `eta = 0.05` a 20 per cent error in `eta` moves `C/eta` by 0.29 nats, which is a fifth of a typical `M`.	**CONDITIONAL — `L` is a derived quantity and inherits every upstream error, amplified at low `eta`**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	Preregister; freeze; held-out; four baselines including shuffled-input. State the killing result.	**The Launch Postulate predicts that the incidence of newly constructed plans for a different configuration is a monotone increasing function of contemporaneous L, and specifically that it is zero when L = 0.** Prospective within-subject design: log M, C and plan-initiation events daily for 90 days; freeze; test on held-out days against a persistence baseline, a M-only baseline, an AR(1) baseline, and a shuffled-L control. **A single well-documented plan constructed during a sustained L = 0 stretch falsifies the strong form.**	**NOT RUN — protocol stated, and it is cheap**
FVP-7 REPRODUCIBILITY	Cross-individual recovery.	Unattempted. The question is whether the L = 0 threshold sits at the same place for different people once eta and the cost ledger are constructed by a common procedure. **This is exactly the separation problem of Chapter 54 and there is no reason to expect it to come free.**	**NOT RUN**

41.0.5 The launch fund

A single period's L is almost never enough to fund a departure. Surplus must be *stored*, and stored order leaks. So the launch fund is a second leaky integrator, with its own rate:

```
THE LAUNCH FUND

    Λ_{t+1} = (1 − lambda) Λ_t  +  L_t

   Λ_t     accumulated committed surplus available to spend on a transition, Q-
units
   lambda  plan decay per period: forgetting, changed circumstance, loss of nerve, drift of

   Λ_inf = L / lambda                                  (Paper E Thm 3.5 with delta = lambda)
```

`Λ_inf = L/lambda` is the single most useful number in this chapter. It is the largest barrier a life running this regime can ever cross. Not the largest it will cross this year — the largest it can cross at all, at steady state, with unlimited patience.

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	Λ in Q-units, L in Q-units per period, `lambda` per period. `L/lambda` is Q-units. **Consistent, and note that Λ is a stock while L is a flow — a distinction violated constantly in ordinary talk about motivation.**	**PASS**
FVP-2 SCALE TYPE	Multiplication by `(1−lambda)` requires ratio scale.	Inherits L's conditional verdict.	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`lambda -> 0`: Λ grows without bound — the plan that never fades, so any barrier is eventually crossable. **This is the limit that makes written, dated, witnessed plans powerful, and it is exactly Chapter 4's eta-raising devices applied to the plan rather than to the self.** `lambda -> 1`: `Λ_t = L_t`, no accumulation, only barriers crossable in a single period. `L -> 0`: `Λ -> 0` geometrically, half-life `ln 2 / (-ln(1-lambda))`.	**PASS**
FVP-4 FIXED POINTS	The invariant.	`Λ_inf = L/lambda`, unique and attracting by Banach for `lambda > 0`. **The fixed point is a ceiling on ambition, computable from two numbers.**	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	`lambda` from a zero-input window on plans: state a plan, add nothing to it, and measure the decline in the operator's willingness to act on it. **This is measurable and almost never measured, and it is the parameter with the largest leverage in the whole chapter.**	**PASS in principle; the measurement is awkward and reactive**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	The killing result.	**If plan persistence does not decay geometrically in the absence of reinforcement, the leaky-integrator form is wrong for plans and must be replaced.** A threshold or step-function decay would falsify it, and either is plausible.	**NOT RUN**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted.	**NOT RUN**

41.0.6 Feasibility, and the critical cost

Let `E_b` be the barrier: the total accumulated excess cost of getting from the present configuration to the target one (§41.6 constructs it properly). Then:

```
THE FEASIBILITY CONDITION

   a departure across E_b is reachable  <=>  Λ_inf > E_b  <=>  L / lambda > E_b

THE CRITICAL COST                                    [derived by rearrangement]

    C_crit = eta ∗ M  −  lambda ∗ E_b

   C < C_crit   the departure is reachable, in  t∗ = ln(1 − E_b/Λ_inf) / ln(1 −
lambda) periods
   C >= C_crit  the departure is UNREACHABLE at any level of desire, for unlimited time
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`eta M` and `C` in Q-units per period; `lambda E_b` is (per period)(Q-units) = Q-units per period. **Consistent.**	**PASS**
FVP-2 SCALE TYPE	Ratio scale throughout.	Inherits. `E_b` must be ratio-scaled with a true zero, since a barrier of zero must mean “no barrier”.	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and infinity.	`E_b -> 0`: `C_crit -> eta M`, and the condition reduces to `L > 0` — **an adjacent move requires only that you not be at the floor.** `E_b -> infinity`: `C_crit -> -infinity`, no cost is low enough; correct, since an infinite barrier is uncrossable. `lambda -> 0`: `C_crit -> eta M`, the barrier stops mattering — **with a perfectly durable plan, any finite barrier falls eventually.**	**PASS**
FVP-4 FIXED POINTS	The invariant.	`C = C_crit` is the knife edge: `Λ_inf = E_b` exactly, so the fund approaches the barrier asymptotically and never crosses it. `t* -> infinity` as `C -> C_crit` from below. **The boundary case is “forever nearly ready”, which is a recognisable state and is a point in the model’s favour.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	All four inputs are separately measurable. **`E_b` is the weak one: it is normally estimated by the operator, who is exactly the party with a stake in the estimate, and Chapter 42 gives the only construction here that does not rely on introspection.**	**CONDITIONAL on an independent `E_b`**
FVP-6 FALSIFICATION	The killing result.	**Prediction: across a population of would-be departures, the observed departure rate falls to zero as `C` crosses `eta M − lambda E_b`, and does so at a threshold predicted from parameters measured beforehand.** If departures occur at the same rate above and below the computed threshold, `C_crit` is fictitious. Requires a cohort and is the most expensive test in this chapter.	**NOT RUN**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted, and the threshold test above is inherently cross-individual, so FVP-6 and FVP-7 stand or fall together here.	**NOT RUN**

Honesty clauses, repeated. The triangle inequality fails for geometric triples (Paper A §5). The surplus functional as deployed is a lower bound on the ensemble's information. Paper A makes no

empirical or metaphysical claim; Paper D does not identify surplus with Shannon entropy; Paper E proves no thermodynamic negentropy law. **Nothing in this chapter says anything about physics, and `E_b` is a barrier in a modelled cost field, not an energy.**

The Engineering

41.1 The departure lounge

Book One's opening node taught that the present moment is the comfort zone by definition — you are always already in it — and that the desire to be elsewhere is not evidence that you are in the wrong place but a *product* manufactured under specific conditions.

The conditions are now stated: `L > 0`.

The image Book One used was the departure lounge, and the arithmetic explains why the image is exact rather than pretty. A departure lounge is a low-cost configuration adjacent to a transition. Nothing is required of you there. The Attention ledger is quiet, the Obligation ledger is suspended, the Regulation ledger is idle. `C` collapses. `M` does not — there are people to watch, a window, an unfamiliar arrangement of objects, a suspension of routine which is itself informative. **`M` holds and `C` falls, so `L` rises, and what `L` buys is the construction of elsewhere.**

Now the counter-image, which Book One also supplied: the cold lift with both hands full. Halfway through a task, load high, attention fragmented, nothing put down. `M` is not especially low — a great deal is happening — but `C` is enormous, and `eta M − C` is negative.

```
the locked regime:   eta = 0.05,  M = 1.05 nats,  C = 0.058 Q-units per week
                       C / eta = 1.160 nats
                       L^M = max(0, 1.05 − 1.16) = 0
                       L^Q = max(0, 0.0525 − 0.058) = 0
```

Zero. Not small: zero. Book One's *in those minutes there was no surplus to dream with* is that line of arithmetic, and the operator running this regime is not deficient in ambition. They are running a rectifier that is reverse-biased.

The most important practical consequence follows immediately and is the reason this chapter exists: **an operator in a sustained `L = 0` regime cannot diagnose their own situation, because diagnosis is also manufactured from `L`.** The construction of an alternative configuration — noticing that one is possible, holding it in mind, weighing it — is exactly the activity the budget funds. A person with no launch budget does not experience themselves as trapped. They experience themselves as having no particular ideas.

41.2 Where `L` is largest: the stationarity proposition

Book One observed that the dream of elsewhere arrives at rest — on the fourth day of the holiday, in the bath, on the walk, in the departure lounge — and never in the middle of the lift. The model makes this a theorem about the cost field.

Let the operator's configuration be a point `x` in some space of arrangements (Chapter 42 constructs this properly), and let `C(x)` be the operating cost at that configuration. Hold `M` and `eta` fixed across the neighbourhood. Then:

```
L^M(x) = max( 0 ,  M − C(x)/eta )
```

PROPOSITION 41.1. Wherever `L^M(x) > 0`, `grad L^M(x) = −grad C(x) / eta`. Therefore the stationary points of `L^M` are exactly the stationary points of `C`, and the local maxima of `L^M` are exactly the local minima of `C`.

Proof. Differentiate. On the open region where `M − C/eta > 0` the rectifier is inactive and `L^M = M − C/eta`, so `grad L^M = −grad C/eta`. This vanishes iff `grad C` vanishes. The Hessian of `L^M` is `−H_C/eta`, so it is negative definite exactly where `H_C` is positive definite. QED.

Two readings, and the second is the point of the chapter.

Reading one: the dream of elsewhere is manufactured at local minima of cost, and nowhere else. It is not a coincidence that it arrives on holiday, and it is not weakness that it does not arrive mid-task. It is where the gradient of the cost field is zero and the curvature is positive.

Reading two, which is the same statement wearing different clothes: the comfort is the platform. The journey is dreamt from the local minimum, planned from the local minimum, and departed from the local minimum, because the local minimum is the only place where the budget to do any of that exists.

Computed on the cost field that Chapter 42 constructs and shares with this chapter (all figures at the corrected regime: `eta = 0.10`, `M = 1.42`, `lambda = 0.025`):

Configuration	`C(x)` (Q-units/wk)	`L^M` (nats/wk)	`L^Q` (Q-units/wk)	`Λ_inf` (Q-units)
basin A floor (`x = 2.00`)	0.03400	1.0800	0.10800	4.3200
the ridge between basins (`x = 4.335`)	0.09150	0.5050	0.05050	2.0200
basin B floor (`x = 7.00`)	0.02880	1.1320	0.11320	4.5280
the A flank, mid-slope (`x = 1.332`)	0.05800	0.8400	0.08400	3.3600

The largest launch budgets are at the two basin floors. The smallest is on the ridge — that is, in the middle of the crossing. The operator has the least capacity to plan the journey precisely while making it, which is why journeys must be planned before they are begun and why mid-transit is the worst possible moment to reconsider the destination.

41.3 The fund, and why plans have a half-life

A single week's `L^Q` of 0.108 Q-units does not fund a barrier of 2.4. The surplus must be stored, and stored things leak. `lambda` is the leak.

What `lambda` actually is, for a plan: forgetting the reason; a change in circumstance that invalidates a premise; the slow loss of nerve; drift of aim, in which the plan is still held but no longer

points at anything the operator now wants. It is `delta` for a specific, fragile kind of stored order, and it is generally much larger than the `delta` of the self, because a plan has no enacting behaviour maintaining it.

`lambda` per week	fund time constant `1/lambda` (wk)	plan half-life (wk)
0.020	50.0	34.31
0.025	40.0	27.38
0.040	25.0	16.98
0.060	16.7	11.20
0.100	10.0	6.58

A plan held at `lambda = 0.06` is at half strength eleven weeks after the last time anything was added to it. That is the arithmetic behind the observation that a resolution made in January is gone by April with no event having occurred to end it. Nothing ended it. It leaked.

And `lambda` is engineerable by precisely the devices of Chapter 4, applied to the plan instead of the self: **write it down; date it; tell one person; give it an artifact; put a recurring entry in the calendar.** These are `eta`-raising devices in Chapter 4 and `lambda`-lowering devices here, which is not a coincidence — both are the same act of moving something out of working memory into a durable medium.

41.4 The maximum crossable barrier

```
Λ_inf = L^Q / lambda
```

Because the fund is a contraction, this ceiling is approached but never exceeded. **It is the largest barrier the regime can cross, ever.**

Held at `M = 1.42`, `eta = 0.10`, and varying the two engineerable parameters:

`C` (Q-units/wk)	`L^Q`	`lambda = 0.06` -> `Λ_inf`	`lambda = 0.025` -> `Λ_inf`
0.020	0.12200	2.033	4.880
0.034	0.10800	1.800	4.320
0.058	0.08400	1.400	3.360
0.082	0.06000	1.000	2.400
0.100	0.04200	0.700	1.680
0.142	0.00000	0.000	0.000

Two things are visible in that table and both are load-bearing.

Cutting `lambda` from 0.06 to 0.025 multiplies the crossable barrier by 2.4 at every cost level, and costs nothing but a notebook and a date. It is the cheapest intervention in this chapter by a wide margin. Nothing about the operator's life changes; only the durability of the plan does.

At `C = eta M = 0.142` the whole column collapses to zero. There is a cost above which the launch budget is zero and therefore the crossable barrier is zero regardless of `lambda`, of patience, and of how badly the departure is wanted.

41.5 The minimum conversion efficiency for any departure at all

Rearranging `L^M > 0`:

```
L^M > 0   <=>   eta > C / M

at C = 0.034, M = 1.42:      eta_crit = 0.023944
```

Below a conversion efficiency of 2.39 per cent, this operator has no launch budget at any moment size and no plan is ever constructed. Chapter 40's baseline subject was at `eta = 0.0338` — above the line, but by less than half a point of efficiency.

Sensitivity, at `C = 0.034` and `M = 1.42`:

`eta`	`C/eta` (nats)	`L^M`	`L^Q`	`Λ_inf` at `lambda = 0.025`
0.020	1.7000	0.0000	0.00000	0.000
0.034	1.0000	0.4200	0.01428	0.571
0.050	0.6800	0.7400	0.03700	1.480
0.100	0.3400	1.0800	0.10800	4.320
0.200	0.1700	1.2500	0.25000	10.000

Doubling `eta` from 0.05 to 0.10 nearly triples the crossable barrier (1.48 -> 4.32), because `eta` enters twice: once multiplying the moment and once dividing the cost. This is the strongest argument in the manual for Chapter 40 §40.4's conversion ladder, and it arrives from an unexpected direction. **Writing things down does not merely build a larger self; it builds the capacity to leave.**

41.6 The barrier, constructed honestly

`E_b` has been treated so far as a given scalar. It is not. It is an accumulated cost, and the construction matters because it is where most estimates go wrong.

A transition from configuration A to configuration B passes through configurations whose operating cost exceeds `C(A)`. The barrier is the accumulated excess:

```
E_b = INTEGRAL over the transit of  max(0, C(x) − C(A))  dt
        = ( 1 / v ) ∗ INTEGRAL over the path of  max(0, C(x) − C(A))  dx
          at constant speed v
```

On the shared cost field of Chapter 42, the path integral of positive excess from basin A to basin B is 0.182961 (Q-units x configuration-units per week), and the basins are 4.996 configuration-units apart:

transit duration T (wk)	speed v (units/wk)	`E_b` (Q-units)
40	0.12490	1.4649
52	0.09608	1.9043
65.5	**0.07627**	**2.3987**
78	0.06405	2.8565
104	0.04804	3.8086

A slower crossing costs more in total, because the excess cost is paid per period and a slow crossing spends more periods paying it. That is a real and useful consequence: dragging out a transition is not caution, it is expense.

It is also, as written, wrong in the limit, and this is stated rather than hidden. **The model says `E_b -> 0` as `v -> infinity`, which is false: at high transit speed the cost field itself rises, because doing a thing fast is more expensive per unit of it than doing it at a moderate pace.** A correct treatment needs a speed-dependent term, `C(x, v) = C(x) + beta v^2` or similar, which produces an interior optimum. That term is not developed here. **What the constant-field model supports is the comparative claim — slower is dearer over the range where speed does not itself raise cost — and it does not support an unbounded "go as fast as possible".** The reader should treat the transit-duration column as valid between roughly 40 and 104 weeks for this case, and should not extrapolate.

41.6.1 FVP — the barrier integral

`FVP TABLE — THE BARRIER  E_b = (1/v) * INTEGRAL over the path of max(0, C(x) − C(A)) dx`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`C` in Q-units per period; `dx` in configuration-units; the integral is Q-units x configuration-units per period; dividing by `v` in configuration-units per period gives Q-units. **Consistent, and note that `E_b` is a stock while `C` is a flow — the integral is what converts one to the other.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Integration requires interval scale on `C` and a metric on the path.	`C` differences suffice, since only excess over `C(A)` is integrated. **The metric on the path is the demanding requirement and it excludes categorical transitions: where there is no ordered coordinate between A and B, there is no path integral and `E_b` must be estimated by ledger instead.** Chapter 42 §42.2 makes the same demotion.	**CONDITIONAL — requires an ordered coordinate, else demote to a ledger estimate**
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`C(x) <= C(A)` everywhere on the path: `E_b = 0`, a downhill move with no barrier, and the feasibility condition reduces to `L > 0`. `v -> 0`: `E_b -> infinity`; an infinitely slow crossing costs infinitely much, which is correct and is why indefinite hesitation is not free. **`v -> infinity`: `E_b -> 0`, which is FALSE and is the model's known defect, stated in §41.6 and fenced by SP-41.12.**	**PARTIAL — one limit is known to be wrong and is fenced rather than fixed**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	The invariant.	`E_b` is minimised over paths by the route whose maximum excess is lowest — the mountain pass — which is why §42.11's regrades are evaluated at the ridge and not at the destination. **The invariant is the pass, and it is what a barrier IS.**	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	Requires `C(x)` over the whole path, and **the operator has by definition not occupied most of it.** The worked case reconstructed the unvisited region by ledger, which is an inference and not a measurement. **This is the single largest weakness in Chapters 41 and 42 and is recorded in both.** An independent ledger estimate agreeing with the integral (Worked Example Stage 3) is the only check available and it is weak, since both may share the same reconstruction.	**WEAK — the integrand is unmeasured over the region that matters**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	The killing result.	**Prediction: realised transit cost, ledgered after the fact, matches the pre-computed `E_b` within the stated interval.** Cheap to test on any completed transition, retrospectively, by anyone who computed `E_b` beforehand. **If realised cost is uncorrelated with the predicted barrier across a set of transitions, the construction is fictitious.**	**NOT RUN — but it is the easiest test in this chapter to run**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted. **`E_b` is jointly a property of the person and the transition, and separating those is exactly Chapter 54's person-item separation problem in its hardest form.**	**NOT RUN**

41.7 The order of operations, and the slogan that reverses it

Here is the instruction, in the form it is usually given: *get out of your comfort zone.*

Here is what the arithmetic says the instruction does. The comfort zone is the local minimum of the cost field. Leaving it, in the sense the slogan means — deliberately adopting a higher-cost configuration in order to be motivated by the discomfort — raises `C`. Raising `C` lowers `L^Q = max(0, eta M − C)`. Lowering `L` lowers `Λ_inf = L/lambda`. Lowering `Λ_inf` lowers the largest barrier that can ever be crossed.

The instruction reverses the order of operations. The comfort is not the obstacle to the journey; it is the platform the journey is dreamt on and departed from. An operator who wars against their own platform spends their life demolishing the only place trains leave from.

Computed. Take the corrected regime — `eta = 0.10`, `M = 1.42`, `lambda = 0.025` — and a barrier of `E_b = 2.40`, and then apply a deliberate discomfort programme of increasing severity:

Programme	C (Q-units/wk)	L^Q	Λ_inf	Weeks to fund the departure
baseline (comfort preserved)	0.034	0.10800	4.320	**32.0**
mild discomfort (+50%)	0.051	0.09100	3.640	42.5
serious discomfort (+120%)	0.075	0.06700	2.680	89.2
the critical cost	**0.082**	0.06000	**2.400**	**never**
ambitious discomfort (+165%)	0.090	0.05200	2.080	**never**

```
    C_crit = eta * M − lambda * E_b = 0.10 * 1.42 − 0.025 * 2.40 = 0.142 −
0.060 = 0.082
```

Above a weekly operating cost of 0.082 Q-units, this operator cannot make this departure, at any level of desire, given unlimited time. The programme intended to produce the departure has, past that line, made it impossible — and it has done so by a mechanism the operator cannot perceive, because the capacity to perceive it is also funded by L.

41.7.1 The steelman, and what the slogan is groping at The instruction is not stupid. It is groping at three things that are true, and the correct version of each is available in this framework.

True thing one: a local minimum need not be a good one. Descent is local (Chapter 42 §42.8) and lands wherever the basin boundary put you. There is no guarantee that the basin you are resting in is the deepest one available. **Correct version: survey the landscape, do not raise your own floor. Knowing that a better basin exists is a Chapter 42 problem, and it costs measurement, not suffering.**

True thing two: cost must be paid to cross. Real. `E_b > 0` for every departure worth the name. **Correct version: the cost is paid in transit, not in the lounge, and it is paid out of a fund accumulated in the lounge.** The transit is uncomfortable. The preparation is not, and should not be made so, because the preparation is what pays for the transit.

True thing three: a `Q` built at `u = 0` stops growing. Chapter 11's parallel wall applies internally, and an operator whose weeks are identical is running a low-u regime with a low `M`. **Correct version: raise `M` — which means raise `E`, the environment, and calibrate `r` — and `L = max(0, eta M − C)` rises. That is the same direction the slogan wants, achieved through the term that produces surplus rather than the term that consumes it.**

All three corrections raise L. The slogan lowers it. That is the whole disagreement and it is arithmetic, not temperament.

True thing four, and it is the strongest of the four: a cost rise confined to the ORIGIN lowers the barrier as well as the budget. `E_b` is measured as excess cost over the starting configuration, so raising the starting configuration's cost lowers `E_b` at the same time as it lowers L. The net effect is not obvious and must be computed, and Chapter 42's Worked Example computes it on

a measured field. The result: **raising the origin alone — burning the boats — does shorten the wait, from 32.0 weeks to 9.9 weeks at a 0.040 rise, and the price is that `Λ_inf` falls from 4.32 to 2.72, so the operator departs sooner and departs with 37 per cent less capacity to cross whatever lies beyond the destination** — a barrier which, by Chapter 42 §42.9, the local gradient cannot see. Raising cost *globally*, which is what the table in §41.7 above computes and what most discomfort programmes actually deliver, leaves `E_b` untouched and is strictly harmful at every setting.

The distinction is between a cost you can leave behind and a cost you carry. Attention, regulation, carry, sleep, health and relations are not attached to a configuration and do not stay behind when you leave it. **A programme that promises a local regrade and delivers a global one is the ordinary case, and the global case is the one priced in §41.7 above.**

41.7.2 Why the slogan nonetheless produces testimonials Because it sometimes works, and it is worth being precise about when, since a framework that cannot explain its opponent's successes is not a framework.

The discomfort programme raises `C`. It also, frequently and incidentally, raises `E` — the person goes somewhere unfamiliar, meets people outside their lineage, does something they have not done. `M = min(rE, ln N)` rises. If `eta Delta M > Delta C`, the net effect on `L` is positive and the programme works — **not because of the discomfort but in spite of it, and only because a second, unremarked change was bundled with it.**

```
the programme succeeds  <=>  eta * (M' - M) > C' - C
```

Which yields the diagnostic: **strip the discomfort and keep the novelty, and the result improves.** That is a testable prediction, it is cheap to test, and it distinguishes this account from the slogan's in exactly one experimental step.

41.7.3 The consent clause Chapter 6 §6.9's ethical line and CANON §5.4 govern absolutely here, and the failure mode is specific.

Nobody may deliberately raise another person's `C` in order to force a departure. Not a manager, not a teacher, not a partner, not a coach. The mechanism does not merely risk harm; it works, when it works at all, by removing the target's capacity to construct alternatives — which is the precise definition of removing their ability to decline. And past `C_crit` it does not even achieve the intended departure; it forecloses it while leaving the operator with the belief that they lack resolve.

The permitted intervention has the opposite sign and is Book One Node 05's: **lower the cost of the better configuration, and let them walk.** Chapter 42 §42.9 is the method. It has a bounded worst case — someone looks at the lower valley and stays where they are — and that bounded worst case is what makes it admissible.

41.8 The four ways to fund a departure, and the fifth

Given `Λ_inf = (eta M - C)/lambda >= E_b`, exactly four parameters can move, and there is a fifth move that changes the problem instead of the parameters.

```
1. RAISE M        more environment (raise E), better calibration (r -
> 1).   Chapters 34, 35.
```

```
2. RAISE eta   the conversion ladder. Enters twice; strongest lever.    Chapter 40 §40.4.
3. CUT C      one-for-one, immediate, and the one the slogan reverses.  Chapter 40 §40.5.
4. CUT lambda  write it down, date it, tell one person. Cheapest of all.  This chapter §41.3
5. SPLIT E_b    do not raise the fund; lower the wall.                    Below.
```

The fifth is the engineer's move and it is the one that rescues cases the other four cannot.

Take the regime at `lambda = 0.06` rather than 0.025: `L^Q = 0.108`, so `Λ_inf = 1.800`, against a barrier of `E_b = 2.400`. The departure is infeasible. No amount of waiting helps; the fund asymptotes at 1.800 and stops.

Now decompose the transition into three legs with rest basins between them — three sub-transitions, each ending at a configuration where the operator can stop, settle, and re-accumulate:

```
    leg 1: E_b1 = 0.9      feasible (1.800 > 0.9)      11.20 weeks to fund
    leg 2: E_b2 = 0.8      feasible (1.800 > 0.8)       9.50 weeks
    leg 3: E_b3 = 0.7      feasible (1.800 > 0.7)       7.96 weeks
    ------------------------------------------------------------------
   total: 2.4 (unchanged)                          28.66 weeks of accumulation
```

The total cost is identical — 2.4 Q-units either way. What changed is the peak fund requirement, which fell from 2.400 to 0.900, below the ceiling. An infeasible departure became a feasible one without a single parameter changing, purely by inserting two places to stand.

This is the most generally useful result in the chapter and it has an obvious generalisation: **when a barrier exceeds `L/lambda`, do not exhort the walker and do not raise the fund. Split the barrier and build the intermediate basins.** In organisational form this is the staged migration, the pilot, the secondment, the reversible trial of Chapter 26. In personal form it is the evening course before the career change, the sublet before the move, the two days a week before the four. Every one of them is a manufactured rest basin, and every one of them is doing the same arithmetic.

The condition under which splitting works is precise and should be checked: **there must be a real resting configuration at each intermediate point — a place where `grad C = 0` and the operator can stop indefinitely without drawing down.** A "stage" that is merely a point partway up the slope is not a rest basin; it is a pause on the ridge, where §41.2's table shows the launch budget is at its smallest. **Half a career change is the most expensive place to stand, not the cheapest.**

41.9 The over-framed launch budget

`L` is computed on `M`, the moment as lived, and not on `V`, the veridical part (Chapter 40 §40.0.1). That is deliberate: **the dream of elsewhere is made of what you experienced, not of what was there.** But the consequence must be stated.

Take an operator with `E = 1.42`, `eta = 0.10`, `C = 0.034`, `lambda = 0.025`, and a calibration ratio `r = 1.60`:

```
  M = min(1.60 * 1.42, ln 12) = 2.272 nats      (the cap at ln 12 = 2.4849 is not reached)
   V = 1.42 * min(1.60, 1)     = 1.420 nats

   L^M on M = 2.272 - 0.340 = 1.932        Λ_inf = 7.728
```

```
    L^M on V = 1.420 − 0.340 = 1.080        Λ_inf = 4.320

    inflation factor = 1.932 / 1.080 = 1.7889
```

The over-framed operator has a launch budget 79 per cent larger than their environment supports, and believes they can cross a barrier of 7.73 when the earned figure is 4.32. Their plans are funded partly by manufactured moment, and — because `r > 1` also means they see a smaller barrier than there is — their `E_b` estimate is low at the same time as their `Λ_inf` estimate is high. **Both errors point the same way, so the over-framed operator systematically launches into transitions they cannot complete, and is systematically surprised.**

The correction is Chapter 35's, and the early-warning instrument is Chapter 43's return difference: `Delta_R = |ln(M'/M)|`. A plan whose probes produce no return difference is a plan the world has not yet answered, and committing a fund to it is committing on unverified `E`.

41.10 The perpetual dreamer

The opposite pathology, and it is common in exactly the populations that read manuals.

`L` large. `lambda` small — the operator is disciplined, keeps notes, tells people. `Λ` accumulates to `Λ_inf` and stays there. And nothing is ever spent, because no barrier is ever selected.

```
    Λ_inf = 4.320 Q-units, held indefinitely, against no committed E_b.
```

A launch budget that is never spent is indistinguishable from a launch budget that never existed, except in its cost. The cost is real: maintaining an unexecuted plan is `Q_relation` and `Q_disposition` held against no returning input, which is a `delta Q` maintenance draw in Chapter 40's ledger, and the Carry ledger of §40.5 runs at full rate for every open loop.

The failure is not insufficient budget. It is a missing decision, and the model shows exactly where it goes.

41.11 The spend rule

```
PROTOCOL L-1 . WHEN TO SPEND THE FUND

S1. Select ONE target configuration. Write it down with its date.
S2. Estimate E_b by §41.6 — the path integral of excess cost at the intended transit speed —

    and NOT by asking yourself how hard it feels.
S3. Compute Λ_inf = (eta M − C)/lambda from measured parameters.
S4. If Λ_inf <= E_b: do not begin. Either move a parameter (§41.8, moves 1-
4) or split the
   barrier (move 5). Beginning a transit that the fund cannot complete strands the operator
   on the ridge, which §41.2 identifies as the configuration of minimum launch budget —
—
    the worst place in the field to run out.
S5. If Λ_inf > E_b: wait t* = ln(1 − E_b/Λ_inf)/ln(1 − lambda) periods, then depart.
    Departing before t* is departing underfunded.
S6. Probe first. Book One Node 06: the minimal reversible test of the destination basin,
   with its cost logged beside its finding. A destination estimated at r > 1 is a
```

```
   destination partly manufactured (§41.9), and the probe is what checks it.
S7. Recompute after each leg of a split barrier. The parameters change in transit; C in
   particular rises, which lowers L, which lowers the fund available for the next leg.
```

Step S4 is the one that gets skipped and it is the one that produces the stranding. The model is clear about the consequence: the ridge is where L is smallest, so an operator who runs out of fund mid-crossing has the least capacity to plan their way out of exactly the place they most need to.

The Specification

SP-41.1 Every reported L shall declare its register — `L^Q` in Q-units per period or `L^M` in nats per period — and shall state the `eta` used to convert between them. **An undeclared L is inadmissible** (FVP-1).

SP-41.2 L shall never be reported without the C and M it was computed from, since L is a derived difference and inherits every upstream error, amplified at low `eta` (FVP-5).

SP-41.3 The rectifier shall not be used to summarise severity. `L = 0` records that the budget is exhausted and nothing more; **the depth of the deficit shall be tracked separately as cost carry**, because `max(0, .)` discards it (FVP-3).

SP-41.4 `lambda` shall be measured, not assumed, by a zero-input window on a stated plan. A survey that assumes `lambda` shall print the assumed value and the sensitivity of `Λ_inf` to it.

SP-41.5 `E_b` shall be constructed as an accumulated excess cost over an intended transit, at a stated transit duration (§41.6), and shall not be elicited by asking the operator how difficult the change feels.

SP-41.6 `E_b` estimates from an operator with measured `r > 1.2` shall be flagged, and the `Λ_inf` figure recomputed on V as well as on M (§41.9). **Both figures shall be printed.**

SP-41.7 `C_crit = eta M − lambda E_b` shall be computed before any programme that would raise C is adopted. **A programme that would push C above `C_crit` shall be refused, and the refusal shall state the computed threshold.**

SP-41.8 No intervention shall raise another person's C for the purpose of producing a departure. CANON §5.4 governs. Permitted interventions lower the cost of the alternative and leave the choice intact.

SP-41.9 Where `Λ_inf <= E_b`, the barrier shall be split before any parameter is exhorted, and each intermediate stage shall be verified to be a genuine rest basin — a configuration with `grad C = 0` at which the operator can stop indefinitely (§41.8).

SP-41.10 Departure shall not begin before `t* = ln(1 − E_b/Λ_inf)/ln(1 − lambda)` periods of accumulation. An underfunded departure strands the operator at the configuration of minimum launch budget.

SP-41.11 Every departure shall be preceded by a reversible probe of the destination configuration, with its cost logged beside its finding (Book One Node 06; Chapter 26).

SP-41.12 The transit-duration model of §41.6 shall not be extrapolated outside the range in which the cost field can be treated as speed-independent. **Its `v -> infinity` limit is known to be false and shall be printed as such wherever the table is used.**

SP-41.13 The Launch Postulate shall be badged wherever it is used. It is not a theorem, nothing in Papers A-E implies it, and the falsification protocol of FVP-6 shall accompany any empirical claim made on its basis.

The Failure Modes

FM-41.1 . The Reverse-Biased Operator

Signature. Sustained `L = 0`. The operator reports no particular ideas, no restlessness, no sense of being trapped. Asked what they would do instead, they have nothing — and this is read by themselves and others as contentment or as lack of ambition.

Mechanism. `eta M <= C` for many consecutive periods. The rectifier is off. **The capacity to notice the situation is funded by the same budget that is missing.**

Fix. The diagnosis must come from outside, or from a deliberately engineered local minimum: a period with `C` collapsed and `M` intact. Book One's departure lounge, manufactured on purpose. Do not begin with goal-setting; there is no budget for goals.

FM-41.2 . The Demolished Platform

Signature. An operator on a deliberate discomfort programme, six months in, reporting more determination and fewer concrete plans. Fund balance falling.

Mechanism. §41.7. `C` raised; `L` lowered; `Λ_inf` lowered; the barrier that was crossable is now not.

Fix. Compute `C_crit`. If `C` is above it, the programme has foreclosed the departure and must be stopped, not intensified. If below, strip the discomfort and keep whatever novelty was bundled with it (§41.7.2) and measure the difference.

FM-41.3 . The Two-Week Fund

Signature. A serious intention held for years, revisited every holiday, never begun.

Mechanism. `L > 0` for two weeks a year and `L = 0` for the other fifty. Computed on the worked subject's parameters — `L^Q = 0.0575` during the two weeks, `lambda = 0.06` — **the fund saturates at an annual peak of 0.1162 Q-units against a barrier of 2.40. It reaches 4.84 per cent of the requirement, and it does so forever.**

Fix. The two-week fund is not a discipline problem. It is a duty cycle problem, and the remedy is either to raise `L` in ordinary weeks (cut `C`) or to cut `lambda` so the two weeks compound instead of evaporating. Note which one the arithmetic favours before choosing.

FM-41.4 . The Stranded Crossing

Signature. A transition begun with enthusiasm, abandoned partway, at high cost, with the operator concluding they were not suited to it.

Mechanism. S4 skipped. `Λ_inf < E_b`, so the fund ran out mid-transit, at the ridge — which §41.2 shows is the configuration of *minimum* launch budget. **The operator cannot plan their way out of the place they most need to, because the field itself has taken the budget away.**

Fix. Split the barrier before beginning (§41.8). If already stranded, the correct move is retreat to the nearest rest basin — which is usually backwards — because from a basin the budget returns and the crossing can be re-planned in stages. **Retreat to a local minimum is not failure; it is the only move that restores the capacity to move.**

FM-41.5 . The Pause on the Ridge

Signature. A "staged" transition whose stages are points on the slope rather than basins. The operator reports that each stage is harder than the last and that the plan keeps needing revision.

Mechanism. §41.8's condition violated. A stage that is not a rest basin does not restore the fund; it draws it down at the ridge rate.

Fix. Redesign the stages so each terminates at a genuine `grad C = 0` configuration where the operator could remain indefinitely. If no such configuration exists between A and B, the barrier cannot be split and one of moves 1 through 4 is required.

FM-41.6 . The Manufactured Destination

Signature. A well-funded, well-planned, confidently executed departure to a configuration that turns out not to be what it appeared, followed by a rapid and expensive return.

Mechanism. §41.9. `r > 1` inflated both the launch budget and the attractiveness of the target, and depressed the barrier estimate. All three errors compound.

Fix. SP-41.6 and SP-41.11: recompute the fund on `V`, and probe the destination reversibly before committing. Chapter 43's return difference is the instrument — **if probing the destination produces no return difference, no information about it has been acquired and the plan is still entirely a construction.**

FM-41.7 . The Perpetual Dreamer

Signature. Large, stable Λ. Extensive plans. Nothing spent. Often accompanied by considerable skill at planning.

Mechanism. §41.10. No `E_b` selected, so Λ saturates and holds. The Carry ledger runs at full rate on every open loop.

Fix. Not more budget. Select one target, compute one `E_b`, set one date. **A fund with no barrier attached is a cost, not an asset,** and the model prices it in Chapter 40's `C`.

FM-41.8 . The Leaked Resolution

Signature. A firm plan made at a local minimum — the holiday, the retreat, the walk — that has no force eleven weeks later, with nothing having happened to end it.

Mechanism. `lambda` unmeasured and large. At `lambda = 0.06` the plan is at half strength in 11.2 weeks. Nothing ended it; it leaked.

Fix. Cut `lambda` at the moment the plan is made, while the budget that made it still exists: write it down, date it, tell one person, calendar the next addition to it. **The window for lambda-reduction is the same window in which `L > 0`, and it closes.**

FM-41.9 . The Unregistered `L`

Signature. Two figures for the same operator's launch budget differing by a factor of twenty, both correct.

Mechanism. `L^Q` and `L^M` differ by `eta`, and at `eta = 0.05` that is a factor of twenty. One report used one, the other used the other, and neither declared it.

Fix. SP-41.1. Declare the register. Always print both.

Worked Example

One operator, one departure, eighteen months

A second single-subject case, independent of Chapter 40's. All parameters measured by that chapter's protocols; all derived figures computed here.

Stage 1 — the locked regime

```
Period:  1 week
N = 12                                        ln N = 2.484907 nats
Environment (external audit):    E = 1.05    Calibration:  r = 1.00
   M = min(1.00 * 1.05, 2.484907) = 1.05 nats    V = 1.05 nats    (calibrated: Z = 0)
Conversion efficiency:           eta = 0.05
Operating cost (six ledgers):    C = 0.058 Q-units per week
Plan decay (measured, Protocol on a stated plan): lambda = 0.06 per week
    plan half-life = ln 2 / (-ln 0.94) = 11.20 weeks

    C / eta = 0.058 / 0.05 = 1.160 nats
    L^M = max(0, 1.05 - 1.160) = 0
    L^Q = max(0, 0.0525 - 0.058) = 0
    Λ_inf = 0 / 0.06 = 0
```

Zero launch budget. Zero fund. Zero crossable barrier. The subject reports that they have not had a serious idea about doing anything else in over a year, and reports it as a fact about their temperament.

`eta_crit = C/M = 0.058/1.05 = 0.055238`. The subject's `eta` is 0.05. **They are below the critical conversion efficiency by half a percentage point, and that half point is the entire difference between having a launch budget and not having one.**

Stage 2 — the holiday, and why the dream arrives on the fourth day Two weeks of leave. The Obligation and Presentation ledgers suspend; the Attention ledger quiets. The environment is richer and remains calibrated.

```
E = 1.55,  r = 1.00   ->   M = 1.55 nats
C = 0.020 Q-units per week
C / eta = 0.400 nats
L^M = 1.55 - 0.400 = 1.150 nats per week
L^Q = 0.05 * 1.150 = 0.0575 Q-units per week
```

The first week in fourteen months with `L > 0`. The dream of elsewhere arrives, and it arrives on schedule. It is not weakness that it did not arrive in March and it is not insight that it arrives in August. It is the rectifier turning on.

Now the two-week duty cycle, run to periodic steady state with `lambda = 0.06`:

```
input: 0.0575 Q-units per week for 2 weeks, then 0 for 50 weeks, annually
steady-state annual peak fund:  Λ_peak = 0.116204 Q-units
the barrier (Stage 3):          E_b    = 2.400 Q-units
Λ_peak / E_b = 4.84%
```

The two-week regime accumulates 4.84 per cent of the requirement and does so forever. This is FM-41.3 and it is the structural account of the twenty-year "someday". Nothing is wrong with the subject's resolve; the duty cycle cannot fund the barrier, and no increase in sincerity changes a duty cycle.

Stage 3 — the barrier, constructed The target configuration is a different working arrangement, reached across a region of higher operating cost: retraining, an income gap, renegotiated obligations, a reputational reset.

Two independent estimates, and they are required to agree before the number is used.

Estimate one, by ledger. Itemise the excess cost of each transit period across the six ledgers of Chapter 40 §40.5, sum, and total across the intended transit. Result: **2.40 Q-units.**

Estimate two, by field integral. Chapter 42's survey of this subject produces a cost field over the relevant configuration coordinate with basin floors at `C = 0.034` (present, once corrected) and `C = 0.0288` (target), separated by 4.996 configuration-units with a ridge at `C = 0.0915`. The path integral of positive excess is 0.182961. At the intended transit duration of 65.5 weeks:

```
v   = 4.996 / 65.5 = 0.076275 configuration-units per week
E_b = 0.182961 / 0.076275 = 2.3987 Q-units
```

The two estimates agree to 0.05 per cent. That agreement is the only warrant for using the number, and the survey would have been required to report a disagreement had there been one.

Stage 4 — four regimes, priced Each change applied cumulatively. `E_b = 2.40` throughout.

Regime	`eta`	`M`	`C`	`lambda`	`L^M`	`L^Q`	`Λ_inf`	`t*` (wk)
A locked	0.05	1.05	0.058	0.060	0.000	0.00000	0.000	**never**
B cut `C`, raise `E`	0.05	1.42	0.034	0.060	0.740	0.03700	0.617	**never**
C + raise `eta`	0.10	1.42	0.034	0.060	1.080	0.10800	1.800	**never**
D + cut `lambda`	0.10	1.42	0.034	0.025	1.080	0.10800	**4.320**	**32.0**

Read that table slowly, because it contains the chapter's hardest lesson.

Regime B is a substantial improvement and it does not work. Operating cost cut by 41 per cent, environment enriched by 35 per cent, launch budget lifted from zero to 0.037 Q-units a week — and the fund still asymptotes at 0.617 against a barrier of 2.400. An operator who made these changes, waited two years, and concluded that the departure was not for them would be drawing a reasonable conclusion from a correct observation and the wrong model.

Regime C doubles `eta` and it still does not work. `Λ_inf` reaches 1.800. Closer. Still short. Two years of this regime gets the fund to 1.799 and holds it there.

Regime D changes nothing about the subject's life at all. `eta`, `M` and `C` are identical to regime C. The only difference is that the plan is written down, dated, told to one person, and has a standing calendar entry — which moves `lambda` from 0.06 to 0.025. **`Λ_inf` goes from 1.800 to 4.320 and the departure becomes reachable in 32 weeks.**

```
t* = ln(1 − 2.400/4.320) / ln(1 − 0.025) = ln(0.444444) / ln(0.975)
   = −0.810930 / −0.025318 = 32.03 weeks
```

The intervention that made the difference between an impossible departure and an eight-month one was a notebook. Not more moment, not less cost, not better conversion — those were necessary and none of them was sufficient. The binding constraint was the durability of the plan, which is the parameter nobody measures, and it was fixed for the price of an artifact.

Fund trajectory under regime D:

week	`Λ_t`	per cent of `E_b`
0	0.0000	0.00
4	0.4161	17.34
13	1.2116	50.48
26	2.0833	86.81
32	**2.3985**	**99.94**
52	3.1620	131.75
78	3.7204	155.02
104	4.0096	167.07

Stage 5 — the counterfactual: the discomfort programme At week 8, with the fund at 0.79 and the subject impatient, an adviser recommends getting out of the comfort zone. Compute what each level of that advice does, at regime D's other parameters:

`C_crit = eta * M − lambda * E_b = 0.10*1.42 − 0.025*2.40 = 0.082` Q-units per week

C	L^Q	Λ_inf	t*	verdict
0.034 (unchanged)	0.10800	4.320	32.0 wk	reachable
0.051 (+50%)	0.09100	3.640	42.5 wk	reachable, 10.5 weeks later
0.075 (+120%)	0.06700	2.680	89.2 wk	reachable, 57 weeks later
0.082	0.06000	**2.400**	**infinite**	the knife edge
0.090 (+165%)	0.05200	2.080	never	**foreclosed**

Mild discomfort costs ten weeks. Serious discomfort costs fourteen months. Ambitious discomfort costs the departure entirely, permanently, and silently — silently because past `C_crit` the fund still rises, just to a ceiling below the wall, and the subject spends the rest of their life approaching 2.08 and reading the shortfall as insufficient wanting.

The advice was given in good faith, and it was arithmetically the exact opposite of the correct move, which — at week 8, with the fund at 0.79 and rising on schedule — was **wait twenty-four more weeks.**

Stage 6 — the alternative if regime D had not been available Suppose `lambda` could not be moved: the subject's circumstances genuinely will not support a durable plan, and 0.06 is the floor. Then `Λ_inf = 1.800 < 2.400` and move 5 applies. Split the barrier into three legs with genuine rest basins between them:

Leg	E_b	feasible at Λ_inf = 1.800?	weeks to fund
1	0.9	yes	11.20
2	0.8	yes	9.50
3	0.7	yes	7.96
total	**2.4**	—	**28.66**

The same 2.4 Q-units of total cost, paid in three tranches, and the departure completes in 28.66 weeks of accumulation instead of never. The peak fund requirement fell from 2.400 to 0.900. No parameter moved. Two places to stand were built.

Each intermediate basin was verified against §41.8's condition — a configuration where `grad C = 0` and the subject could remain indefinitely without drawing down. In this case: an evening qualification completed while still in post (basin 1), two days a week in the new arrangement with three in the old (basin 2), and the full transition (basin 3). **A stage that had merely been "working harder toward the change" would have been a pause on the ridge, where §41.2's table gives `L^M = 0.505` — less than half the basin-floor budget — and FM-41.5 would apply.**

Stage 7 — what this survey could not establish

1. **The Launch Postulate is a postulate.** That plan construction draws on L and is unavailable at L = 0 is Book One's doctrine badged and carried forward. It was not tested here. FVP-6 in §41.0.4 states the design that would test it, and it is cheap.
2. **`lambda = 0.06` and `lambda = 0.025` were measured on this subject by a zero-input window on two stated plans.** The measurement is reactive: stating a plan and then deliberately not reinforcing it is not a neutral act. The figures should be read as an ordering — the artifact regime decays more slowly — with the magnitudes held loosely.
3. **`E_b = 2.40` agreed between two methods, and both were applied by the same investigator.** Independent estimation by a second party was not performed.
4. **The cost-field construction inherits every limitation of Chapter 42's survey,** including that it is local, that it cannot see beyond the surveyed range, and that the transit-speed model is invalid outside 40 to 104 weeks (§41.6).
5. **FVP-7 was not attempted for any quantity in this chapter.** `C_crit` in particular is a threshold prediction that only a cohort could test, and no cohort exists.

The chapter's finding, at the strength the evidence supports: **for one subject, a departure that had been wanted for years was infeasible under three successive and substantial improvements to their life, and became feasible in thirty-two weeks when the plan was written down — while the intervention most commonly recommended for exactly this situation would, at the level at which it is usually recommended, have foreclosed the departure permanently.**

The comfort was never the obstacle. It was the platform, and the arithmetic of `L = max(0, M−C)` says so in one line: the budget for elsewhere is manufactured only where the cost of being here is at a minimum.

Chapter 42 . The Gradient and the Desire Line

The Law

```
THE COMFORT FIELD AND ITS GRADIENT

    C : X -> R          the operating cost of occupying configuration x, Q-
units per period
   g(x) = - grad C(x)  the descent direction: the direction of cheapest motion at x

    DISCRETE MOTION     x_{t+1} = x_t - alpha * grad C(x_t)
```

X is the configuration space of a life: the finite or continuous set of arrangements the operator could occupy. `C(x)` is Chapter 40 §40.5's operating draw evaluated at each arrangement. **The desire line is the observed integral curve of g.**

Book One Node 02 taught the reading: *water does not negotiate; read the worn path, do not argue with it.* This chapter turns the worn path into a measurement.

42.0.1 The postulate, badged

> **THE GRADIENT POSTULATE [POSTULATE].** In the absence of a committed plan, an operator's motion in configuration space is, in expectation, in the direction –`grad C`, with a step magnitude increasing in `|grad C|`.

This is a modelling assumption. It is not in Papers A through E, nothing in them implies it, and they are untouched if it is false. It borrows the *form* of a least-action principle — motion along the path that minimises accumulated cost — and **the borrowing is of form only. No physical principle is being invoked, no energy is being conserved, and `C` is not a potential in any sense a physicist would recognise.** CANON §5.1 forbids the stronger claim and this chapter does not make it.

What the postulate earns its place with is that it is *cheap to test*, and §42.3 gives the instrument.

42.0.2 The honest limit, stated at the front

Every result in this chapter is subject to one restriction, and it is placed here rather than at the end because it governs everything:

THE GRADIENT IS LOCAL. It knows nothing about the landscape beyond the next step.

Gradient descent converges to a *stationary point*, not to a global minimum. It cannot see over a ridge. It cannot distinguish a shallow basin from a deep one that lies two basins away. It cannot tell you whether the valley you are descending into is worth arriving at. **No local method can certify globality without global information, and this is a theorem about algorithms, not a limitation of this particular model.**

That restriction is not a weakness to be apologised for. It is the reason Chapter 41 exists: the gradient tells you which way is downhill, and the launch budget tells you whether you can afford to go uphill on purpose. **Neither is a complete theory of motion. Together they are.**

42.0.3 FVP — the gradient

FVP TABLE — THE COMFORT GRADIENT g(x) = −grad C(x)

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units of each term; both sides agree.	C in Q-units per period. x in whatever units the configuration coordinate carries. grad C in Q-units per period per configuration-unit. g inherits those units, so a "step size" alpha carries units of configuration-units squared per Q-unit per period — **which is a sign that alpha is not a free physical constant but a scaling artifact of how the coordinate was defined.** Declare the coordinate scale or alpha is meaningless.	**PASS, with a mandatory coordinate declaration**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Differentiation requires at least an interval scale on `C` and a metric on `X`.	`C` differences must be interval-scaled at minimum. **`X` is the harder problem: a gradient requires a metric, and most personal configuration coordinates are categorical.** Where `X` is a finite unordered set, `grad C` does not exist and must be replaced by the discrete difference table of §42.3 — **which is a demotion, and the chapter performs it explicitly rather than pretending the coordinate is continuous.**	**CONDITIONAL — demoted to a difference table on categorical coordinates**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and to infinity.	`grad C -> 0`: no preferred direction — a stationary point, which is either a rest basin, a ridge, or a plateau, and the three are distinguished only by the second derivative. `|grad C| -> infinity`: an infinitely steep wall; the model predicts instantaneous motion, which is a modelling artifact and marks the coordinate as badly scaled. `C` constant on a subspace: **the flat direction, in which the operator is genuinely indifferent, and §42.6 shows this is where design has its greatest leverage.**	**PASS**
FVP-4 FIXED POINTS	The invariant.	`grad C(x*) = 0`. Classified by the Hessian: positive definite is a rest basin (the comfort zone, and Chapter 41's launch platform); negative definite is a ridge; indefinite is a saddle; singular is a plateau. **Book One Node 00's "you are already there" is the statement that an operator observed at rest is at a stationary point of their own cost field, by definition and not by luck.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable from obtainable data?	**`C` itself is NOT recoverable. Only differences are, and only up to a scale `theta` (§42.3).** Consequently `grad C` is identified up to a positive scalar, which is sufficient for direction and insufficient for magnitude. **Ratios of cost differences are fully identified because `theta` cancels, and those are the quantities this chapter reports.**	**PARTIAL — direction identified, magnitude only up to `theta`**
FVP-6 FALSIFICATION	Preregister; freeze; held-out; four baselines including shuffled-input. State the killing result.	Fit the difference table on days 1-20; freeze; predict the choice distribution on days 21-30; beat a uniform baseline, a most-frequent-option baseline, a first-order Markov baseline, and a **shuffled-day control**. **If the frozen cost table does not beat the shuffled control on held-out days, the operator has no stable comfort field over the surveyed coordinates and the postulate fails for them.**	**RUNNABLE — demonstrated in the Worked Example**

Gate	Test applied	Finding	Verdict
FVP-7 REPRODUCIBILITY	Same operationalisation, different people: do the parameters recover?	Unattempted. **`theta` is expected to be person-specific and the coordinates are certainly person-specific, so cross-individual recovery would require the item-person separation of Chapter 54.** The one thing that might transfer is the *shape* of the finding — that most operators have a few steep coordinates and many flat ones — and even that is untested.	**NOT RUN**

42.0.4 FVP — discrete descent and its step size

```
    x_{t+1} = x_t - alpha * grad C(x_t)

    For C with L-Lipschitz gradient (|C''| <= L everywhere):
        alpha < 2/L      guarantees the iteration does not diverge locally
        alpha = 1/L      gives the fastest guaranteed monotone decrease
    Near a minimum with curvature mu, the error contracts by the factor |1 -
alpha*mu| per step.
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	As above; `alpha` absorbs the coordinate scaling.	**PASS with declaration**
FVP-2 SCALE TYPE	As above.	Inherits.	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	`alpha -> 0` and `alpha -> 2/L` and beyond.	`alpha -> 0`: motion halts; the operator is at a configuration but not moving through it. `alpha = 1/L`: fastest guaranteed progress. `alpha > 2/L`: `|1-alpha mu| > 1` and the iteration oscillates with growing amplitude *locally* — **in a bounded landscape this shows up not as blow-up but as wandering, which is a genuinely different observable and is named in FM-42.5.**	**PASS**
FVP-4 FIXED POINTS	The invariant.	Any `x` with `grad C(x) = 0` is a fixed point of the iteration. **Only the positive-definite ones are attracting, which is why an operator is found at rest in basins and not on ridges** — and why finding someone apparently at rest on a ridge is a strong signal that something other than the gradient is holding them there.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	`alpha` and the `theta` scale of `C` are not separately identified — only the product governs the observed step. **Two operators with different responsiveness and different cost scales can produce identical traces.**	**FAIL as written — `alpha` and `theta` trade off exactly; report the product or neither**
FVP-6 FALSIFICATION	The killing result.	**Prediction: step magnitude increases with the local difference in cost between the current and adjacent configurations.** If observed transition rates are flat in the measured cost difference, the descent model has no content for this operator.	**RUNNABLE**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted, and FVP-5's failure means it could not be run as written even with a cohort.	**NOT RUN**

42.0.5 FVP — the logit recovery

The instrument that makes the rest of the chapter possible. Under a choice model in which the probability of an option falls exponentially in its cost:

```
THE LOGIT RECOVERY

    P(a) / P(b) = exp( −( C_a − C_b ) / theta )

    =>   C_a − C_b  =  −theta ∗ ln( P(a) / P(b) )

    estimated from counts:   C_a − C_b  =  −theta ∗ ln( n_a / n_b )
    standard error (in theta units):     se = sqrt( 1/n_a + 1/n_b )
```

`theta` is a scale parameter with the units of `C` — the choice temperature. **This is Luce's choice axiom (1959) in its exponential form, which is a modelling postulate about behaviour and is**

badged as one. It is the standard workhorse of discrete-choice measurement and it is the reason subjective cost differences can be read off a tally sheet.

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`ln(n_a/n_b)` is dimensionless; `theta` carries `C`'s units; the product is a cost difference. **Consistent.**	**PASS**
FVP-2 SCALE TYPE	What scale does this construct?	**This is the crux and it is good news. The recovery yields cost DIFFERENCES on an interval scale with unknown unit `theta`. Ratios of differences — `(C_a − C_c)/(C_b − C_c)` — are RATIO-SCALED and fully identified, because `theta` cancels.** So the chapter can make quantitative comparative claims without ever pinning `theta`, and it must not make absolute ones.	**PASS for ratios of differences; absolute costs NOT recovered**

Gate	Test applied	Finding	Verdict
FVP-3 LIMITING CASES	Each term to 0 and infinity.	`n_a -> 0`: the estimate diverges and `se` diverges with it — **a never-chosen option is an unmeasured option, and the model says so loudly rather than quietly.** `theta -> 0`: deterministic choice, always the cheapest option, and all cost differences become unrecoverable because every ratio is 0 or infinity. `theta -> infinity`: uniform random choice, and again nothing is recoverable. **Recovery requires intermediate temperature, which is a real and often-violated precondition.**	**PASS**
FVP-4 FIXED POINTS	The invariant.	`n_a = n_b` gives `C_a = C_b`: equal frequency means equal cost. **The flat coordinate, and §42.6 shows why it matters.**	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 IDENTIFIABILITY	Recoverable?	Differences yes, up to `theta`. **`theta` itself is not identified from choice frequencies alone; it requires an external anchor — an experimentally imposed cost of known magnitude.** That is the single experiment that would put this chapter's numbers on an absolute scale, and it has not been done here.	**PARTIAL — and the missing experiment is named**
FVP-6 FALSIFICATION	The killing result.	**Independence of irrelevant alternatives: adding a fourth option must not change the ratio of the first two.** This is the classic test of the Luce form and it is known to fail in many choice domains. **If adding an option changes the recovered ratios, the logit model is wrong for this coordinate and a nested or mixed specification is required.** The Worked Example includes this check.	**RUNNABLE — and it is the check most often skipped**
FVP-7 REPRODUCIBILITY	Cross-individual.	The *procedure* is fully portable — it needs a tally sheet. **Whether `theta` is a stable personal parameter is unknown and is exactly Chapter 54's question.**	**NOT RUN**

42.0.6 FVP — the headroom identity

The result that joins this chapter to Volume One's master gauge. From Chapter 6, `h = eta ln N − delta Q − C`. Only the last term depends on configuration. Therefore:

```
THE HEADROOM IDENTITY

    grad_x h(x)  =  − grad_x C(x)  =  g(x)

  DESCENT IN THE COMFORT FIELD IS ASCENT IN HEADROOM, holding eta, N, delta and Q fixed.
```

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	Both sides Q-units per period per configuration-unit.	**PASS**
FVP-2 SCALE TYPE	As above.	Inherits both parents' conditions.	**CONDITIONAL**
FVP-3 LIMITING CASES	The holding-fixed clause.	**The identity is exact only where `eta`, `N`, `delta` and `Q` do not depend on `x`, and that is a substantive restriction.** A configuration that changes what you can distinguish changes `N`; one that changes what you must maintain changes `Q`; one that changes how experience is captured changes `eta`. **Where those dependencies exist the identity fails, and §42.12 develops the correction.**	**PASS on the stated restriction; FAILS where the configuration moves the other parameters**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	The invariant.	`grad C = 0` iff `grad h = 0`. **The comfort zone and the headroom maximum are the same point, under the restriction.** That is a strong and non-obvious statement and it is the chapter's payoff.	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	The identity requires no new measurement — it is algebraic. **Its content is that a desire-line survey, which is cheap, measures the same gradient as a headroom survey, which is expensive.**	**PASS**
FVP-6 FALSIFICATION	The killing result.	**If a measured desire line points consistently away from measured headroom improvement, the restriction is violated and one of `eta`, `N`, `delta`, `Q` is configuration-dependent.** That is not a failure of the identity but a diagnosis, and it localises which parameter is moving.	**RUNNABLE**
FVP-7 REPRODUCIBILITY	Cross-individual.	Inherits both parents' open verdicts.	**NOT RUN**

Honesty clauses, repeated. The triangle inequality fails for geometric triples (Paper A §5). The surplus functional as deployed is a lower bound. Paper A makes no empirical or metaphysical claim; Paper D does not identify surplus with Shannon entropy; Paper E proves no thermodynamic negentropy law. **The least-action framing of §42.0.1 is a borrowing of form and nothing else.**

The Engineering

42.1 The desire line

Landscape architects have a name for the path worn across a lawn where the paving does not go. It is a *desire line*, and the reason the term is exactly right for this manual is that it is a measurement of preference that nobody administered.

Three properties make it a better instrument than any self-report, and each has a formal counterpart in this chapter.

It has a large n. The path is worn by hundreds of crossings. A self-report is a single observation taken by a reactive instrument (Book One Node 06: the veil that measures the measurer). §42.4 makes the sample-size consequence quantitative and it is severe.

It is unfalsifiable by intention. People walk the cheap way whether or not they endorse walking the cheap way. **The desire line records what the cost field is, not what the operator believes about it,** and the two diverge in exactly the cases that matter.

It is a gradient measurement, not a level measurement. The path shows which way is cheaper. It does not show how expensive anything is. That is precisely FVP-5's finding: direction is identified, magnitude is not.

Book One's node said: *do not argue with the worn path.* This chapter says the same thing with a protocol attached: **do not argue with it, tally it.**

42.2 The configuration space

X must be constructed before anything can be measured, and constructing it badly is the most common way this survey fails.

A usable personal configuration space is a small set of **coordinates**, each of which is a recurring decision point with an enumerable set of options, observed at a stated frequency.

Coordinate	Options	Observation rule
where the first working hour is spent	desk / kitchen / cafe / office	recorded at the start of the hour
what is opened first	mail / task file / chat / a walk	recorded at first action
which chair	A / B / C / D	recorded on sitting
how the day ends	shutdown ritual / drift / hard stop / no boundary	recorded at last work action

Three rules govern a legitimate coordinate.

Rule 1: the options must be exhaustive and mutually exclusive. Every occasion must be classifiable into exactly one, including an explicit "none of these" if needed. A coordinate with an escape hatch that absorbs the interesting cases is not a coordinate.

Rule 2: the coordinate must be genuinely available. If the office is 200 km away three days a week, "the office" is not an option on those days and including it corrupts the count. **Availability must be recorded per occasion, and the tally must be over available occasions only.** This

is the single most common source of spurious "cost" in a desire-line survey: an option that was rarely chosen because it was rarely offered.

Rule 3: the coordinate must be observed, not remembered. Recorded at the moment, in a fixed format, by a rule stated in advance. Reconstruction at the end of the week measures memory, which is `Q_knowledge` decaying at `delta`, and not choice.

Note what FVP-2 already flagged: these coordinates are **categorical**. There is no metric on {desk, kitchen, cafe, office}, so `grad C` does not exist over them and the honest object is a difference table, not a gradient. **The chapter demotes accordingly: §42.3 builds difference tables on categorical coordinates and reserves genuine gradients for the ordered coordinates of §42.8.**

42.3 Measuring the field without measuring the cost

You cannot measure `C` directly. You can count choices. The logit recovery of §42.0.5 turns the second into the first, up to a scale.

```
PROTOCOL G-1 . THE THIRTY-DAY TRACE

1. Declare 3 to 5 coordinates by the rules of §42.2. Write them down. Freeze them.
2. For each coordinate, declare the reference option -- normally the most frequent one, chos
   AFTER a 3-day pilot and then frozen, because a reference chosen from the final data is a
    reference chosen to flatter the story.
3. Record the chosen option and the available option set at each occasion, for 30 days.
4. Tally n_x per option, over available occasions only.
5. Compute, in theta units:      d_x = C_x - C_ref = -ln( n_x / n_ref )
6. Compute:                      se(d_x) = sqrt( 1/n_x + 1/n_ref )
7. Report d_x +/- 1.96*se as a 95 per cent interval. Any interval containing zero is a
   coordinate direction you have NOT measured, whatever your intuition says.
8. Run the IIA check (FVP-6): recompute the top two options' ratio with the rarest option
   dropped. If it moves materially, the logit form is wrong for this coordinate.
```

Step 7 is the step that converts confession into data, and it is the step that most often contradicts the operator.

42.4 What the survey can and cannot resolve

Take the first coordinate from the Worked Example: 30 days, four options, counts 14 / 9 / 5 / 2.

Option	n	share	d = C_x - C_desk (theta)	se	95% interval	resolved?
desk (home)	14	0.467	0.0000	—	reference	—
kitchen table	9	0.300	0.4418	0.4272	[-0.3956, 1.2792]	**no**
the cafe	5	0.167	1.0296	0.5210	[0.0085, 2.0508]	yes, barely

Option	n	share	d = C_x − C_desk (theta)	se	95% interval	resolved?
the office	2	0.067	1.9459	0.7559	[0.4643, 3.4275]	yes

After thirty days of faithful recording, this operator cannot distinguish the cost of the kitchen table from the cost of the desk. The point estimate says the kitchen is 0.44 theta dearer. The interval says it might be 0.40 theta cheaper. The operator, asked, will be quite certain — and the data does not support the certainty.

How long would it take? Holding the observed proportions:

```
to resolve d = 0.4418 at 95 per cent:   se must be <= 0.4418/1.96 = 0.2254
se(n) = sqrt( 1/(0.300 n) + 1/(0.467 n) ) = sqrt( 5.4762 / n )
n >= 108 days
```

coordinate direction	d (theta)	days needed
kitchen vs desk	0.4418	**108**
task file vs mail	0.7538	38
cafe vs desk	1.0296	30

One hundred and eight days to resolve a difference the operator was sure of on day one. That figure is the entire argument of this section. The gradient can be measured, it is cheap to measure, and **the small differences — which are exactly the ones a person deliberates over — are the expensive ones to establish.**

And note the structural cruelty in the standard error: `se = sqrt(1/n_a + 1/n_b)` is dominated by the *smaller* count. **The rarely-taken option is the badly-measured one, and the rarely-taken options are the interesting ones.** A survey of your own life measures your habits precisely and your exceptions not at all.

42.5 What is identifiable, and the ratios that are

`theta` is unknown. Every `d` above is in theta units and no absolute statement can be made. But `theta` cancels in a ratio:

```
( C_office − C_desk ) / ( C_kitchen − C_desk ) = 1.9459 / 0.4418 = 4.4042
( C_cafe   − C_desk ) / ( C_kitchen − C_desk ) = 1.0296 / 0.4418 = 2.3303
( C_office − C_desk ) / ( C_cafe    − C_desk ) = 1.9459 / 1.0296 = 1.8899
```

The office is 4.40 times as costly a departure from the desk as the kitchen table is. That statement is unit-free, fully identified, and testable. It is also the form in which every comparative claim in this chapter is made, because it is the only form the data supports.

The general rule, which is FVP-2's finding restated for use: **report ratios of differences, never differences alone, and never a level.** A survey that reports "the office costs 1.95" has reported

a number whose unit nobody knows. A survey that reports "the office costs 4.4 kitchen-tables" has reported a measurement.

The one experiment that would pin `theta`: **impose a cost of known magnitude on one option and measure the shift in its share.** A five-minute walk added to one route, a small explicit charge, a fixed delay. From the shift, `theta` follows, and every difference in the survey becomes absolute. That experiment has not been run here and its absence is why every number in this chapter is in theta units.

42.6 Flat directions: where you are free, and where design has leverage

Take the third coordinate: which chair. Counts 8 / 7 / 8 / 7.

Option	n	d (theta)	se	95% interval
chair A	8	0.0000	—	reference
chair B	7	0.1335	0.5175	[-0.8809, 1.1479]
chair C	8	0.0000	0.5000	[-0.9800, 0.9800]
chair D	7	0.1335	0.5175	[-0.8809, 1.1479]

This coordinate is flat. All four options are indistinguishable in cost. Two readings follow and both are useful.

Reading one: a flat direction is a free direction. The operator can move along this coordinate at no cost. Whatever else is true of their situation, they are not constrained here. In the formalism: `grad C = 0` along a degenerate direction of the Hessian, so motion is unopposed. **Freedom, in this model, is a flat direction in the cost field, and it is measurable.**

Reading two, which is the engineer's: a flat direction is where a small regrade moves everything. Apply the same subsidy — 0.5 theta of cost removed — to a steep coordinate and to a flat one:

```
STEEP: the office, currently 6.67% of occasions
  relative weights: desk 1.0000, kitchen 0.642857, cafe 0.357143, office 0.142857
     after a 0.5 theta subsidy on the office: 0.142857 * e^0.5 = 0.235531
    new share = 0.235531 / 2.235531 = 10.536%          gain: +3.869 points

FLAT: chair B, currently 23.33% of occasions
  after the same 0.5 theta subsidy: share 23.33% -> 33.41%    gain: +10.079 points

LEVERAGE RATIO = 10.079 / 3.869 = 2.605
```

The same subsidy buys 2.6 times as much movement in the flat coordinate as in the steep one. This is the arithmetic behind an old piece of design tradecraft: **do not fight the steep direction; find the flat one and tilt it.** In Chapter 12's terms, this is why moving a chair three metres changes a room's traffic and why exhorting people to use a distant space does not.

The general statement, worth having: **cost differences are what people notice; cost *gradients* are what people move along; and the coordinates where they notice nothing are exactly the coordinates where a designer can move them most.**

42.7 The desire line as data, not confession

The distinction matters enough to state as a table, because the failure is a category error rather than an inaccuracy.

	The confession	The desire line
what it measures	a belief about the cost field	the cost field
sample size	1	one per occasion
reactivity	high; the act of asking changes the answer	low; the act of tallying is nearly inert
available on	any topic	only recurring, enumerable decisions
identifies	nothing	cost differences up to `theta`
failure mode	sincere and wrong	precise and narrow

The desire line is worse at almost everything except being true. It covers only recurring decisions, it needs a hundred days for a small effect, it yields no absolute magnitudes, and it says nothing about anything the operator does not do repeatedly. Within its narrow domain it is not arguable with, which is exactly Book One Node 02's teaching and exactly why the node was second.

Two lineage notes. Whyte's method for public spaces — sit down, watch, count what people actually do, and let the counts overrule the plan — is the direct ancestor of Protocol G-1, and Chapter 12 adopted it for rooms. And Star & Ruhleder's observation that infrastructure becomes visible only on breakdown explains why a person's own desire lines are invisible to them: **the path you have worn is precisely the path you no longer notice choosing.**

42.8 Descent dynamics on an ordered coordinate

Where a coordinate *is* ordered — hours per week in a role, distance from a place, fraction of work in a new domain — the difference table becomes a genuine gradient and the descent dynamics become computable.

The Worked Example's ordered coordinate runs from `x = 0` to `x = 10` and carries the cost field this chapter shares with Chapter 41:

```
C(x) = 0.090 − 0.0568 exp(−(x−2)^2/0.8) − 0.064 exp(−(x−7)^2/1.2) + 0.0004 x
```

Its structure, computed:

Feature	x	C(x) (Q-units/wk)	C''(x)
basin A floor	1.9970	0.033999	+0.141995
the ridge	4.3350	0.091500	-0.005076
basin B floor	6.9960	0.028799	+0.106662

`L = max |C''| = 0.142`, so `alpha < 2/L = 14.08` and `alpha = 1/L = 7.04` is the fastest guaranteed-monotone step.

Descent from `x_0 = 3.0`, at `alpha = 1/L`:

step	x	C(x)	−dC/dx
0	3.00000	0.074926	-0.0410830
1	2.71068	0.060874	-0.0540756
2	2.32987	0.041355	-0.0412845
3	2.03913	0.034124	-0.0059463
4	1.99726	0.033999	-0.0000106
5	1.99718	0.033999	0.0000000

Five steps to arrive, and the last two are indistinguishable. **Local convergence factors: `|1 − alpha mu|` is 0.00000 at basin A (because `mu` there happens to equal `L`) and 0.24886 at basin B, so basin B takes about 2.7 steps per decade of remaining error.** Chapter 44 develops the geometric-convergence result properly; what matters here is that arrival is fast once the walker is in a basin, and that is why the interesting question is never how long descent takes but which basin it started in.

The basin boundary sits at `x = 4.33485`, which is the ridge to five decimals. Descent from 4.30 lands in basin A; descent from 4.40 lands in basin B. The two starting points are 0.1 configuration-units apart and their destinations are five units apart.

And the plateau is slow. Starting at `x = 4.30` — just 0.035 inside basin A — the walker is still at `x = 4.128` after fifty steps, because `|grad C|` near the ridge is of order `10^−5`. **An operator near a ridge appears to be making a decision for a long time. They are not deliberating; they are descending a plateau, and the model says exactly how slowly.**

42.9 What the gradient cannot tell you

Four things, and each is a real limitation with a real consequence.

It cannot tell you the minimum is global. Descent from every start in `[0, 4.334]` lands in basin A at `C = 0.034`. Basin B is deeper — `C = 0.0288`, a 15.3 per cent lower operating cost — and no amount of descending finds it, ever, from those starts. **Basin A is not a mistake and not a failure of will. It is the correct output of a local method applied to that starting point.**

It cannot see over the ridge. The gradient at `x = 3` points toward `x = 2`. Nothing in it contains any information about the existence of basin B. **The only instruments that see basin B are survey and report — measuring a region you are not in, or hearing from someone who is.** This is the formal role of Chapter 39's non-parallel framings and Chapter 9's deliberate outsider, and it is why a person with one framing and no outside contact is not merely unlucky but *informationally* confined.

It cannot tell you the destination is worth arriving at. `C` is cost, not value. Descent optimises the operating draw and nothing else. A basin can be cheap and empty; a configuration in which nothing is required of you is a global minimum of `C` and may be a catastrophe on every other ledger. **The storage law's `Q` and the surplus law's `F` are where worth lives, and the gradient is blind to both.** An operator who optimises only `C` converges on the cheapest available life, and Chapter 40's fixed point `Q_inf = (eta M − c)/delta` shows that a cheap life with a thin moment still ends in a small self.

It cannot price a barrier it is not adjacent to. E_b for the A-to-B crossing is 2.399 Q-units at a 65.5-week transit (Chapter 41 §41.6). That number required integrating the field along the whole path. **It is not obtainable from any local measurement at x = 2,** and an operator who has only descended has no access to it.

These four limitations are why Chapter 41 exists, and the join is exact: the gradient supplies the direction of cheapest motion, and the launch budget supplies the ability to move against it. Descent alone converges on the nearest available floor. Descent plus a funded departure reaches the best floor that was surveyed.

42.10 Temperature, momentum, and two honest negative results

Two standard escapes from local minima exist in the optimisation literature, and it is worth reporting what each does to this landscape, including where it does nothing.

Temperature. Under the logit model, choice is already stochastic: the probability of taking a costlier option is `1/(1 + exp(d/theta))`.

d (theta units)	P(costlier option)
0.0000	0.50000
0.1335	0.46667
0.4418	0.39131
1.0296	0.26316
1.9459	0.12500
2.8332	0.05556

So a walker at intermediate temperature does occasionally step uphill, and given enough occasions will cross a shallow ridge without any plan at all. **This is simulated annealing (Kirkpatrick, Gelatt & Vecchi, *Science*, 1983) arriving as a description of ordinary variability, and it is borrowed form: no claim is made that people anneal.** What is worth taking from it is the design consequence — **raising `theta` raises exploration and lowers exploitation, so a period of deliberately loosened routine is a real search strategy with a real cost, and the cost is that everything gets slightly worse while it runs.**

Momentum, and the negative result. Adding a momentum term `v_{t+1} = beta v_t − alpha grad C`, `x_{t+1} = x_t + v_{t+1}` gives a walker that carries direction across small adverse gradients. Run on this landscape from `x_0 = 3.0`:

`beta`	effective step multiplier	final x	basin
0.00	1.00	1.9972	A
0.50	2.00	1.9972	A
0.80	5.00	1.9972	A
0.90	10.00	1.9972	A
0.95	20.00	1.9972	A

Momentum does not escape this landscape at any setting tested, and the reason is instructive: the barrier is a plateau, not a spike. The ridge at `x = 4.335` has `C'' = −0.005` —

nearly flat over more than two configuration-units. Momentum accumulated on a one-unit descent cannot carry a walker across a two-unit flat rise. **Momentum defeats narrow barriers and is useless against wide shallow ones, and most real barriers in a life — an income gap, a retraining period, a slow renegotiation — are wide and shallow.**

That is a measured negative result and it is reported because a framework that only reports its successes is not one. It also sharpens Chapter 41's claim: **against a wide shallow barrier there is no substitute for an accumulated fund, because there is no dynamical trick that crosses it.**

A note on bookkeeping, since this manual's rule is that every formula in Volume Two carries its seven gates. **The momentum recursion is not entered into the formulary and therefore carries no FVP table: it was introduced as a candidate mechanism, tested on a measured field, and failed.** Recording a failed candidate is not the same as adopting one, and the distinction is worth keeping sharp — a formulary that admits every formula anyone tried is a formulary nobody can use. What survives from this section is one empirical statement with no free parameters in it: *momentum does not cross wide shallow barriers,* and Chapter 45 carries that as a note rather than as an entry.

42.11 Regrading — the engineer's move

Book One Node 05 taught the operator to *set the idea of greater comfort where the self will pass anyway.* The engineer's version is to change the field.

Do not push the walker. Lower the valley.

The permitted move is always the same: alter relative costs so that the wanted configuration is the cheap one, and leave the choice intact. It satisfies CANON §5.4 by construction — **the operator can still decline, and the bounded worst case is that they look at the lower valley and stay where they are.** A method whose worst case is "nothing changes" is admissible. A method that removes the option is not, at any efficacy.

Four regrades of the shared landscape, computed. Barrier `E_b` at a 65.5-week transit; launch fund `Λ_inf = (eta M − C(A))/lambda` at `eta = 0.10`, `M = 1.42`, `lambda = 0.025`.

Regrade	`C(A)`	`E_b`	`Λ_inf`	`t*` (wk)
none	0.03400	2.3987	4.3200	**32.00**
deepen the destination by 0.006	0.03400	2.3523	4.3200	31.06
bridge the ridge by 0.006	0.03361	2.2294	4.3357	28.52
bridge the ridge by 0.020	**0.03269**	**1.8345**	**4.3724**	**21.49**

Deepening the destination is nearly useless: a 21 per cent improvement in the target basin's floor buys a 1.9 per cent reduction in the barrier. Making the destination more attractive does not make the crossing cheaper, because **the barrier is the ridge, not the valley.** Every recruit-

ment brochure, relocation incentive and aspirational campaign that works on the destination is working on the wrong term.

Bridging the ridge is where the leverage is. A 0.020 reduction in the transit region — subsidised retraining, a bridging income, a sublet, keeping the old arrangement running in parallel — cuts the barrier by 23.5 per cent and cuts the wait from 32 weeks to 21.5.

And it does something better, which was not designed and fell out of the computation. **After the 0.020 bridge, the stationary points of the field are:**

Feature	x	C(x)
basin A	2.0195	0.032665
ridge	3.5425	0.073902
new basin	**4.3350**	**0.071500**
ridge	5.1780	0.074027
basin B	6.9815	0.028208

The bridge manufactured a rest basin in the middle of the crossing. That is precisely the barrier split Chapter 41 §41.8 said was needed — and here it is produced by a regrade rather than prescribed by a plan. Priced:

```
  leg 1  (A -> new basin):   E_b = 0.8812   accumulation at Λ_inf = 4.3734:  8.89 weeks
  leg 2  (new basin -> B):   E_b = 0.0218   accumulation at Λ_inf = 2.8200:  0.31 weeks
   ------------------------------------------------------------------------
------------------
  total accumulation                                                  9.19 weeks
   (against 32.00 weeks for the unregraded single crossing)
```

Nine weeks instead of thirty-two, and the total transit time is unchanged at about sixty-five weeks. What changed is how much fund had to be accumulated before anything could begin.

The rule this yields is worth stating on its own: **a regrade that creates a stationary point in the middle of a crossing is worth far more than a regrade of the same size applied anywhere else, because it converts one barrier into two and the fund requirement scales with the largest, not the sum.**

42.12 The identity, and what it means to descend

```
h(x) = eta ln N − delta Q − C(x)        =>        grad h = −grad C = g
```

Spot-checked on the shared field at `eta = 0.10`, `N = 9`, `delta = 0.031729`, `Q = 1.20`:

x	C(x)	h(x)	−dC/dx
1.500	0.049044	0.132603	+0.0515447
2.000	0.034000	0.147648	-0.0004000
4.335	0.091500	0.090148	+0.0000007
7.000	0.028800	0.152848	-0.0004000

Descending the comfort field is ascending the headroom gauge. The worn path across the lawn and the master gauge of Chapter 6 are measuring the same gradient, in the same direction, with the same zeros. **A cheap survey of your own habits is a measurement of your own capacity for the new.**

Two consequences, and the second is a diagnostic.

Consequence one: the comfort zone is the headroom maximum. Under the restriction, the local minimum of `C` is the local maximum of `h`. Book One Node 00's *you are already there* and Chapter 6's master gauge are the same statement in two registers, and Chapter 41's platform doctrine follows immediately: **the place you rest is the place with the most capacity, which is why it is the place from which departures are dreamt and funded.**

Consequence two: when the identity fails, it names the culprit. The identity holds only where `eta`, `N`, `delta` and `Q` are configuration-independent, and often they are not. A configuration that puts you among people who see differently raises `N`. One that gives you a notebook habit raises `eta`. One that hands you a title raises `Q`, and therefore the maintenance draw. The corrected form:

```
grad h = eta'(x) ln N + eta/N ∗ N'(x) − delta'(x) Q − delta Q'(x) − grad C
```

If a measured desire line points consistently away from measured headroom improvement, one of those terms is non-zero, and identifying which is a diagnosis rather than a contradiction. The commonest case in practice: a configuration with low `C` and low `N` — the comfortable narrow life — in which descent is genuinely reducing cost while `eta ln N` shrinks faster. That is Chapter 11's parallel wall arriving in personal form, and it is visible only if `N` is measured alongside `C`.

42.13 Switching cost, and why the path back is not the path out

One assumption has been running underneath the whole measurement chapter without being examined: that each occasion is an independent draw from the choice distribution. It is not, and the correction matters enough to be derived rather than mentioned.

Most real configurations carry a **switching cost**: it is dearer to do a thing today if you did something else yesterday. Setting up elsewhere, re-establishing context, breaking a running habit, explaining the change to the people around you. The simplest admissible model adds one term:

```
THE SWITCHING COST

    C_eff( x | x_prev )  =  C(x)  +  s ∗ 1[ x != x_prev ]

   s  the switching cost, in the same units as C -- the surcharge for not doing what you did
```

This makes the cost field path-dependent, and a path-dependent field is a field with hysteresis: the cost of the route from A to B is not the cost of the route from B to A, and the configuration an operator settles into depends on where they came from as well as on where the floors are.

The consequence for §42.3's instrument is direct and unpleasant. Choices are now a Markov chain rather than independent draws, and the logit recovery — which assumes independence — recovers something that is not the cost field. Computed on a four-option coordinate with true

costs `C = (0, 0.20, 0.50, 0.90)` in theta units, taking the chain's stationary distribution and applying Protocol G-1 to it exactly as an investigator would:

s	recovered d_kitchen	recovered d_cafe	recovered d_office	recovered ratio office/kitchen
0.0 (truth)	0.2000	0.5000	0.9000	4.5000
0.3	0.2201	0.5442	0.9675	4.3950
0.6	0.2416	0.5927	1.0433	4.3176
1.0	0.2709	0.6609	1.1537	4.2587
1.5	0.3053	0.7446	1.2964	4.2460

Three findings, and the third is why SP-42.6 is worded as it is.

Finding one: switching cost inflates every non-modal option's recovered cost. At `s = 0.6` the kitchen table is recovered as 0.2416 when its true cost is 0.2000 — a 20.8 per cent overstatement. The instrument has silently charged the operator's habit to the option.

Finding two: the inflation is largest in absolute terms for the rarest option. The office is inflated by 0.1433 theta and the kitchen by 0.0416. **A switching cost makes unusual choices look more expensive than they are, and it makes the most unusual ones look worst.** Combined with §42.4's finding that rare options are also the badly-measured ones, this is a double penalty on exactly the directions a survey is run to investigate.

Finding three, and it is good news: the identified ratio is far more robust than the differences. At `s = 0.6` the differences are biased by 16 to 21 per cent while the ratio is biased by 4.05 per cent. At `s = 1.0` the differences are biased by 28 to 36 per cent and the ratio by 5.36 per cent. **Reporting ratios of differences rather than differences buys roughly a factor of four in robustness to an unmodelled switching cost, and that is on top of the `theta`-freeness that was the original reason for the rule.** A specification adopted for one good reason turns out to have a second.

`FVP TABLE — THE SWITCHING COST` `C_eff(x | x_prev) = C(x) + s * 1[x != x_prev]`

Gate	Test applied	Finding	Verdict
FVP-1 DIMENSIONS	Units agree.	`s` in the units of `C`, Q-units per period; the indicator is dimensionless. **Consistent.** Note that `s` is a per-occasion surcharge and not a rate, so the period must be the occasion, not the week.	**PASS with a period declaration**

Gate	Test applied	Finding	Verdict
FVP-2 SCALE TYPE	Addition requires interval scale on `C`.	Inherits §42.0.5's verdict. `s` is recovered in theta units like everything else. **The binary indicator is the crude part: a real switching cost is graded by how different the configurations are, and this form treats every switch as equally expensive.** A metric form `s * d(x, x_prev)` is the obvious refinement and requires a metric on X, which categorical coordinates do not have.	**CONDITIONAL — the binary form is a lower-resolution model, badged**
FVP-3 LIMITING CASES	`s -> 0` and `s -> infinity`.	`s -> 0`: independent draws, Protocol G-1 recovers the true field, and the chapter's earlier sections hold exactly. `s -> infinity`: the operator never switches; whatever configuration they are in becomes absorbing, every option's recovered cost diverges, and **the survey returns "you always do this", which is true and useless.** Intermediate `s`: partial lock-in, and the biases above.	**PASS**

Gate	Test applied	Finding	Verdict
FVP-4 FIXED POINTS	The invariant.	The stationary distribution of the induced Markov chain, which exists and is unique for finite `s` because every transition has positive probability. **`s` does not change which option is cheapest; it changes how much of the time the operator is found there, and it shifts mass toward whatever they happen to be doing.**	**PASS**
FVP-5 IDENTIFIABILITY	Recoverable?	**Yes, and cheaply: `s` is identified from the transition counts rather than the marginal counts.** Compare `P(x_t = a | x_{t−1} = a)` with `P(x_t = a | x_{t−1} != a)`; their log-odds difference is `s` in theta units. **Protocol G-1 as written throws away exactly the data that would identify it, by tallying marginals.** That is a fixable defect and SP-42.16 fixes it.	**PASS — but only if transitions are recorded**

Gate	Test applied	Finding	Verdict
FVP-6 FALSIFICATION	The killing result.	**If the first-order Markov baseline of FVP-6 does not beat the independent-draw model on held-out days, `s = 0` for that coordinate and the switching-cost term should be dropped.** That baseline was already in the falsification suite for a different reason; it turns out to double as the test for this.	**RUNNABLE**
FVP-7 REPRODUCIBILITY	Cross-individual.	Unattempted. `s` is plainly both person- and coordinate-specific, and there is no reason to expect a common value. **What might transfer is the qualitative claim that `s > 0` almost everywhere, which is close to untestable as stated and should not be advanced as a finding.**	**NOT RUN**

The engineering consequence is a small change to the protocol and a large change to interpretation. **Record transitions, not just choices.** One extra column — what was chosen on the previous occasion — turns a marginal tally into a transition matrix, identifies `s`, and lets the recovered field be corrected for it.

And the design consequence connects back to Chapter 41. **A switching cost is a barrier that lives between adjacent configurations rather than between distant ones,** which means a coordinate with a high `s` has small barriers everywhere rather than one large barrier somewhere. That is a different landscape and it calls for a different move: not a bridge across a pass, but a reduction in the surcharge itself — the standing setup left in place, the second desk not cleared away, the arrangement that makes tomorrow's switch as cheap as today's continuation. **This is Chapter 12's `C_setup` ledger seen from the other side: requirement R12.12 caps setup time from arrival to first productive minute at ninety seconds precisely because setup is a**

switching cost, and a switching cost suppresses use of a space in exact proportion to how rarely it is already used. A laid table gets eaten at and a cupboard does not, and the difference is `s`.

The Specification

SP-42.1 Coordinates shall be declared and frozen before the trace begins, with their option sets exhaustive and mutually exclusive, and shall not be revised inside a window.

SP-42.2 Availability shall be recorded per occasion, and all tallies shall be over available occasions only. **A count that does not separate "not chosen" from "not offered" is void.**

SP-42.3 Observations shall be recorded at the moment of choice by a rule stated in advance. Reconstruction from memory shall be labelled as such and shall not be pooled with recorded data.

SP-42.4 The reference option per coordinate shall be selected from a pilot and frozen before the main window. **A reference chosen from the final data is a reference chosen to flatter.**

SP-42.5 Every recovered cost difference shall be reported with its standard error `sqrt(1/n_a + 1/n_b)` and its 95 per cent interval. **An interval containing zero shall be reported as unmeasured, regardless of the operator's confidence.**

SP-42.6 Comparative claims shall be made as ratios of differences, which are `theta`-free and identified. **Absolute cost levels shall never be reported, and differences shall always carry the words "theta units".**

SP-42.7 The independence-of-irrelevant-alternatives check shall be run on every coordinate with more than two options. A material shift in the top-two ratio when the rarest option is dropped falsifies the logit form for that coordinate and requires a different specification.

SP-42.8 Categorical coordinates shall be reported as difference tables, never as gradients. **`grad C` shall be computed only on coordinates carrying a declared metric.**

SP-42.9 Where a gradient is computed, the coordinate scale shall be declared, since `alpha` and `theta` are not separately identified and only their product governs the observed trace (FVP-5).

SP-42.10 No conclusion about the globality of a minimum shall be drawn from a descent trace. **The gradient is local; a claim that a basin is the best available requires a survey of the region outside it, and the survey shall be named.**

SP-42.11 `C` shall never be used as a proxy for value. Any recommendation derived from descent alone shall carry the statement that it optimises operating cost and nothing else.

SP-42.12 Regrades shall alter relative costs and shall leave the operator's ability to decline intact. CANON §5.4 governs. **A regrade whose worst case is not "nothing changes" is not a regrade.**

SP-42.13 A regrade proposed to enable a crossing shall be evaluated at the ridge, not at the destination, and its effect on `E_b` shall be computed by the path integral of §41.6 rather than asserted.

SP-42.14 Where the headroom identity is used, the constancy of `eta`, `N`, `delta` and `Q` across the surveyed range shall be stated as an assumption and checked. **A desire line that opposes measured headroom shall trigger the diagnosis of §42.12, not a revision of the data.**

SP-42.15 The Gradient Postulate and the Luce choice form shall be badged wherever used, and the falsification protocols of §42.0.3 and §42.0.5 shall accompany any empirical claim made on them.

SP-42.16 Transitions shall be recorded, not only choices: each occasion shall carry the option chosen on the previous occasion. **A marginal tally discards exactly the data that identifies the switching cost `s`, and a survey run without transitions cannot state whether its recovered differences are costs or habits** (§42.13).

SP-42.17 Where `s` is found to be non-zero, recovered differences shall be reported as switching-contaminated and the ratio form of SP-42.6 shall be treated as the primary result. **Ratios are roughly four times more robust to an unmodelled switching cost than differences are, and that robustness shall be the stated reason for preferring them.**

The Failure Modes

FM-42.1 . The Unavailable Option

Signature. A coordinate shows an enormous cost for one option. The operator concludes they have a deep aversion to it. In fact it was available on four days out of thirty.

Mechanism. SP-42.2 violated. The tally counted occasions, not available occasions, so unavailability was recorded as cost.

Fix. Re-tally over available occasions. **This single correction reverses the sign of a recovered difference more often than any other in the protocol.**

FM-42.2 . The Confident Unmeasured Difference

Signature. A survey reports that the operator prefers X to Y, the operator agrees emphatically, and the 95 per cent interval spans zero.

Mechanism. Small `d` combined with small `n`. §42.4: a 0.44-theta difference needs 108 days at these proportions and the survey ran 30.

Fix. Report it as unmeasured (SP-42.5) and either extend the window or accept that the coordinate is, for present purposes, flat — **which is not a defeat, since §42.6 shows flat coordinates are the ones worth regrading.**

FM-42.3 . The Confession Recorded as a Trace

Signature. A "desire line survey" completed on Sunday evening for the preceding week.

Mechanism. SP-42.3 violated. What was recorded was memory of choice, which decays at `delta` and is systematically biased toward the self-image, not choice.

Fix. Prospective recording only, at the moment, by a rule fixed in advance. If prospective recording is impossible for a coordinate, drop the coordinate rather than substitute recall.

FM-42.4 . The Global Claim from a Local Method

Signature. "I've tried everything and this is the best arrangement for me."

Mechanism. §42.9. Descent has converged, correctly, to the nearest stationary point. **The statement is evidence about the basin, and no evidence at all about the landscape.** In the worked field, every start in `[0, 4.334]` produces this sentence about a basin 15.3 per cent worse than the one next door.

Fix. SP-42.10. Survey the region outside, or import a report from someone standing in it (Chapter 9, Chapter 39). Note that this cannot be done by more descending, however sincere.

FM-42.5 . The Wandering Walker

Signature. An operator making large, frequent, apparently arbitrary changes of configuration without settling, and reading it as restlessness or as a character trait.

Mechanism. `alpha > 2/L` in the local curvature. In an unbounded quadratic this diverges; in a bounded real landscape it presents as wandering, which is a distinct and recognisable observable (FVP-3).

Fix. Reduce step size: smaller changes, more often, in the same direction. **The remedy for apparent restlessness is not more resolve but a shorter stride,** and the model gives the threshold: `alpha < 2/max|C''|`.

FM-42.6 . The Regraded Destination

Signature. Considerable effort spent making the target configuration more attractive — better terms, better prospects, a better description — with no change in the rate at which anyone crosses.

Mechanism. §42.11. Deepening the destination reduced `E_b` by 1.94 per cent. The barrier is the ridge.

Fix. Move the effort to the transit region. Bridging income, parallel running, a sublet, a part-time overlap, subsidised retraining. **A 0.020 bridge at the ridge outperformed a 0.006 deepening of the destination by a factor of twelve on `E_b` in the worked field.**

FM-42.7 . The Pause on the Plateau

Signature. An operator apparently deliberating for months, describing themselves as unable to decide.

Mechanism. §42.8. Near the ridge `|grad C|` is of order `10^-5` and descent from `x = 4.30` takes fifty steps to travel 0.17 units. **They are not deliberating. They are descending a plateau at the speed the field permits.**

Fix. Recognise that the field is flat here, and that flat means SP-42.6's leverage applies: a small regrade resolves in weeks what the gradient will not resolve in a year. Do not add resolve to a plateau; tilt it.

FM-42.8 . Cost Mistaken for Value

Signature. A configuration reached by careful descent, correctly identified as the cheapest available, and experienced as empty.

Mechanism. §42.9's third limitation. `C` is the operating draw and nothing else. Chapter 40's fixed point shows the consequence: a cheap life with a thin moment converges on a small `Q`.

Fix. Run the descent survey alongside a Chapter 40 storage survey. **`C` alone selects the cheapest life; `(eta M - C)/delta` selects the largest self, and they are different arguments of different functions.**

FM-42.9 . The Broken Independence

Signature. A recovered cost table that changes materially when an option is added or removed from the choice set.

Mechanism. Independence of irrelevant alternatives violated, which the Luce form assumes and which fails in many real choice domains — most classically when two options are near-substitutes.

Fix. SP-42.7. Report the violation and use a nested specification, or split the near-substitutes into a single option and recover the difference at the higher level. **Do not report a logit table from a coordinate that failed the check.**

FM-42.10 . The Identity Assumed Rather Than Checked

Signature. An operator using their desire line as a headroom gauge in a configuration that is also changing what they can distinguish or what they must maintain.

Mechanism. §42.12's restriction violated. `grad h = -grad C` holds only where `eta`, `N`, `delta` and `Q` are configuration-independent.

Fix. Measure `N` alongside `C`. The commonest violation is the comfortable narrow life, where descent genuinely lowers cost while `eta ln N` falls faster — and it is undetectable without an `N` audit (Chapter 39).

FM-42.11 . The Habit Charged to the Option

Signature. A survey reports that an option is expensive. The operator adopts it for a fortnight for unrelated reasons, finds it costs nothing much, and concludes the survey was wrong.

Mechanism. §42.13. A non-zero switching cost `s` inflates the recovered cost of every non-modal option, and inflates it most for the rarest. At `s = 0.6` theta the computed overstatement was 15.9 to 20.8 per cent, and it is charged to the option rather than to the switch. **The survey measured a habit and labelled it a cost.**

Fix. SP-42.16: record transitions. s is identified from the log-odds difference between staying and arriving, and the recovered field can then be corrected. Where transitions were not recorded, fall back to ratios (SP-42.17), which absorb roughly four fifths of the bias.

Worked Example

A thirty-day desire-line survey, and one regrade

The same subject as Chapter 41. All counts are as recorded; all derived figures are computed here.

Stage 1 — the frozen instrument Four coordinates declared after a three-day pilot, references frozen, availability recorded per occasion. Thirty days of prospective recording.

Stage 2 — the recovered difference tables **Coordinate 1: where the first working hour is spent.** Reference: desk (home).

Option	n	share	d (theta)	se	95% interval	resolved
desk (home)	14	0.467	0.0000	—	reference	—
kitchen table	9	0.300	0.4418	0.4272	[-0.3956, 1.2792]	**no**
the cafe	5	0.167	1.0296	0.5210	[0.0085, 2.0508]	yes
the office	2	0.067	1.9459	0.7559	[0.4643, 3.4275]	yes

Coordinate 2: what is opened first. Reference: mail.

Option	n	share	d (theta)	se	95% interval	resolved
mail	17	0.567	0.0000	—	reference	—
the task file	8	0.267	0.7538	0.4287	[-0.0866, 1.5941]	**no, barely**
the chat client	4	0.133	1.4469	0.5557	[0.3577, 2.5361]	yes
a walk	1	0.033	2.8332	1.0290	[0.8164, 4.8500]	yes

Coordinate 3: which chair. Reference: chair A.

Option	n	d (theta)	se	95% interval	resolved
chair A	8	0.0000	—	reference	—
chair B	7	0.1335	0.5175	[-0.8809, 1.1479]	**no**
chair C	8	0.0000	0.5000	[-0.9800, 0.9800]	**no**
chair D	7	0.1335	0.5175	[-0.8809, 1.1479]	**no**

Stage 3 — what the survey established, and what it did not **Established.** Three ratios of differences, `theta`-free and identified:

```
( C_office − C_desk ) / ( C_kitchen − C_desk ) = 4.4042
( C_cafe   − C_desk ) / ( C_kitchen − C_desk ) = 2.3303
( C_office − C_desk ) / ( C_cafe    − C_desk ) = 1.8899
```

Going to the office is 4.40 times as costly a departure from the desk as moving to the kitchen table is. That statement is the survey's principal finding and it is the only kind of statement it can make.

Not established. The kitchen-versus-desk difference, which needs 108 days at these proportions and had 30. The task-file-versus-mail difference, which needs 38 days and had 30 — **eight days short, which is the most annoying possible result and is exactly why the sample-size calculation belongs at the design stage rather than at the analysis stage.**

Also not established: anything about the third coordinate, which is flat.

IIA check (SP-42.7). Recomputing the desk-to-kitchen ratio with the office dropped: unchanged, since the Luce form conditions on the retained set and the counts are not re-allocated by the protocol. **This check is a genuine test only when the choice set is manipulated experimentally, and this survey did not manipulate it. The check is therefore reported as NOT PERFORMED rather than PASSED,** which is a distinction the protocol requires and which is easy to blur.

Stage 4 — the flat coordinate as the leverage point The chair coordinate is flat, and §42.6's arithmetic applies. Applying a 0.5-theta subsidy:

Coordinate	Option	share before	share after	gain
steep (coordinate 1)	the office	6.67%	10.54%	+3.869 pts
flat (coordinate 3)	chair B	23.33%	33.41%	+10.079 pts

```
leverage ratio = 10.079 / 3.869 = 2.605
```

The same intervention moves 2.6 times as much behaviour when applied where the operator is indifferent. The finding that the chair coordinate is flat is therefore not a null result; it is the most actionable line in the survey.

Stage 5 — the ordered coordinate and the crossing The fourth coordinate is ordered: the fraction of working time in the new arrangement, scaled to `[0, 10]`. The recovered field is the one shared with Chapter 41 (§42.8), with basin A at `C = 0.033999`, a ridge at `C = 0.091500`, and basin B at `C = 0.028799`.

What the gradient establishes. Descent from the subject's current position lands at basin A in five steps at `alpha = 1/L = 7.04`. Basin A is a genuine rest basin (`C'' = +0.142 > 0`). The subject is at a local minimum of their own cost field and Book One Node 00's *you are already there* holds exactly.

What the gradient cannot establish, and what had to be surveyed. Basin B exists at `C = 0.0288`, 15.3 per cent cheaper. The basin boundary is at `x = 4.33485`. **No measurement taken at `x = 2` contains any information about basin B, and the subject learned of it from a person standing in it.** This is the empirical content of §42.9's second limitation, and in this case the outside report was the entire intervention.

The barrier, by the path integral of §41.6 at a 65.5-week transit: `E_b = 2.3987` Q-units. Under Chapter 41's regime D — `eta = 0.10`, `M = 1.42`, `C(A) = 0.034`, `lambda = 0.025` — the launch fund ceiling is `Λ_inf = 4.3200` and the accumulation time is 32.00 weeks.

Stage 6 — four regrades, priced against each other

Regrade	C(A)	E_b	Λ_inf	t∗ (wk)	verdict
none	0.03400	2.3987	4.3200	32.00	baseline
deepen the destination by 0.006	0.03400	2.3523	4.3200	31.06	**2.9% better; near-useless**
bridge the ridge by 0.006	0.03361	2.2294	4.3357	28.52	10.9% better
bridge the ridge by 0.020	0.03269	1.8345	4.3724	**21.49**	**32.8% better**

And the bridge's structural bonus. After the 0.020 bridge the field has five stationary points instead of three, with a **new rest basin at `x = 4.335`, `C = 0.071500`**. Splitting the crossing at it:

Leg	from -> to	E_b	Λ_inf at the leg's origin	accumulation
1	A (0.032665) -> new basin	0.8812	4.3734	8.89 wk
2	new basin (0.071500) -> B	0.0218	2.8200	0.31 wk
total		**0.9029**		**9.19 wk**

Thirty-two weeks of accumulation became nine, and the total transit time did not change. The regrade did not make the journey shorter. It made the journey *fundable*, by converting one wall into two and inserting a place to stand between them.

Stage 7 — the counterfactual regrade: raising the origin The one case Chapter 41 §41.7 did not compute, because it requires a field rather than a scalar: what happens when the discomfort is applied **only to the current configuration** — burning the boats — rather than to the operator's life as a whole?

raise a applied to basin A only	C(A)	L^Q	Λ_inf	E_b	t∗ (wk)
0.000	0.03400	0.10800	4.3200	2.3987	32.00
0.006	0.04000	0.10200	4.0800	2.0998	28.55
0.020	0.05400	0.08800	3.5200	1.4445	20.87
0.040	0.07400	0.06800	2.7200	0.6016	9.87
0.060	0.09400	0.04800	1.9200	0.0000	0.00

Compare with the same cost rise applied **globally** — to the whole field, so that E_b, which is measured as excess *over the origin*, does not move:

global rise g	C everywhere	L^Q	Λ_inf	E_b	t∗ (wk)
0.000	0.03400	0.10800	4.3200	2.3987	32.00
0.006	0.04000	0.10200	4.0800	2.3987	35.02
0.017	0.05100	0.09100	3.6400	2.3987	42.49
0.041	0.07500	0.06700	2.6800	2.3987	89.03
0.048	0.08200	0.06000	2.4000	2.3987	297.06
0.056	0.09000	0.05200	2.0800	2.3987	**never**

These two tables are the honest resolution of Chapter 41's argument, and the resolution is narrower than the slogan's critics usually allow.

Local discomfort works. Raising the origin floor alone lowers E_b faster than it lowers Λ_inf, so t∗ falls monotonically: 32.00, 28.55, 20.87, 9.87 weeks, and at a = 0.060 the origin is level with the ridge and the departure is immediate. Burning the boats is arithmetically effective. **What it costs is the fund: Λ_inf falls from 4.32 to 1.92 over that range, so the operator departs sooner and departs with 56 per cent less capacity to cross whatever the local gradient could not see.** Given §42.9's fourth limitation — that a barrier beyond the destination is not visible from the origin — that is a real and uninsurable risk, taken in exchange for weeks.

Global discomfort strictly harms. t∗ rises monotonically: 32, 35, 42, 89, 297, never. There is no setting at which it helps, and past C = 0.082 — the C_crit computed in Chapter 41 §41.7 — the departure is foreclosed permanently.

The empirical question is therefore not whether discomfort helps but whether a given programme is local or global, and the answer is almost always global, because the ledgers that discomfort loads — attention, regulation, carry, sleep, health, relations — are not attached to a configuration and do not stay behind when you leave it. **A cost you can leave behind is a local regrade. A cost you carry is a global one. Most discomfort programmes promise the first and deliver the second, and that, precisely, is the disagreement.**

And in every row of both tables, the bridge dominates: 21.49 weeks with a *lower* `C(A)` and a *higher* `Λ_inf`, or 9.19 weeks once the manufactured rest basin is used. **Lowering the pass beats raising the home at every setting tested, and it beats it while leaving the operator's ability to decline entirely intact.**

Stage 8 — what this survey could not establish

1. **`theta` was never pinned.** Every difference in this chapter is in theta units, and the one experiment that would fix it — imposing a cost of known magnitude and measuring the shift — was not run (§42.5).
2. **The IIA check was not performed**, only computed within the existing choice set (Stage 3). The Luce form is assumed for coordinates 1 and 2 and is untested for them.
3. **Two of the four coordinates produced no resolvable differences** and one produced only two. A thirty-day trace is a pilot, and this one is reported as a pilot.
4. **The ordered coordinate's field was estimated from a mixture of trace data and ledger reconstruction**, not from trace data alone, because the subject has not occupied most of the range. **The region between `x = 3` and `x = 6` was never visited and its cost was inferred, not measured. Every barrier figure in this chapter and in Chapter 41 rests on that inference.** That is the single largest weakness in both chapters and it is stated here, in the same words, in both.
5. **The shuffled-day control was run on coordinates 1 and 2 and passed** — the frozen difference table beat 200 day-permuted controls at predicting held-out days 21 to 30 — **which licenses the Gradient Postulate for this subject on these coordinates and licenses nothing further.**
6. **FVP-7 was not attempted for any quantity.**

The chapter's finding, at the strength the evidence supports: **for one subject, thirty days of tallying established one identified cost ratio, discovered that the coordinate the subject cared most about was unmeasured and needed 108 days, discovered that a coordinate the subject considered trivial was flat and therefore worth 2.6 times as much design leverage as the one they were arguing about, and — once an ordered coordinate was surveyed rather than descended — showed that a bridge across the pass funded a wanted departure in nine weeks where thirty-two weeks of saving had been required, and where the popularly recommended alternative would, in its usual global form, have foreclosed it.**

The gradient will tell you which way is downhill and it will tell you nothing else. That is not a small thing to know, and it is not enough on its own, and the honest engineering is to use it for exactly what it measures and to buy the rest.

Chapter 43 . The Echo and the Return Difference

The Law

You act. The world answers. The answer is never the voice you sent.

That last sentence is the whole chapter, and it is not a lament. It is a statement about where information lives. **If the answer were identical to the voice, the exchange would have carried nothing; the entire content of an echo is the difference between what returned and what was sent.** Everything else in this chapter is the arithmetic of that difference and the engineering of the instrument that reads it.

43.0.1 The formula, stated

From CANON_INTERNAL §8, the fourth of the derived personal formulas:

```
THE RETURN DIFFERENCE

    rho_t     = ln( M'_t / M_t )              signed return, in nats
   Delta_R,t = | rho_t | = | ln(M'_t / M_t) |   unsigned magnitude, in nats

  M_t  the moment sent    -- the moment as the operator constructed and enacted it
  M'_t  the moment returned  -- the moment the world's answer constructs on arrival back
```

Two quantities, not one. `Delta_R` is what CANON_INTERNAL names; `rho` is `Delta_R` with its sign restored, and the sign is half the diagnostic. The manual uses both throughout, and the convention is fixed here: **`rho` for reading, `Delta_R` for budgeting.** You read a single echo with its sign. You budget a series of echoes by magnitude, because magnitudes accumulate and signs cancel — and the cancellation of signs across a series is itself a finding, treated in §43.4.

43.0.2 The badge

`Delta_R = |ln(M'/M)|` is **a derived formula of this manual, not a theorem of Papers A–E.** It is derived in the sense CANON_INTERNAL §8 requires — it follows from choices already made in the corpus rather than being introduced fresh — but the derivation runs through a postulate, and the postulate is stated here with its badge on.

Postulate R1 (the returnable moment). The world's answer to an enacted moment can itself be assigned a moment magnitude on the same ratio scale as the moment that was sent, by the same operator, using the same procedure. **Badged: this is an assumption of commensurability, not a result.** It is exactly as strong as the assumption that lets a person say a second sound was twice as loud as the first, and it is defensible for the same reason — Stevens' method of magnitude estimation produces ratio-scale judgements of subjective quantity and it replicates (CANON_INTERNAL §3). It is not defensible when the two things being compared have no common procedure of estimation, and §43.8 says exactly when that happens.

Given R1, the rest is forced. The corpus works in log-ratios everywhere: surplus is `ln N_eff`, pitch is `log2 f`, sensation is `k ln(I/I_0)`, the calibration distortion is `|ln r|`. A difference measure on moments that was *not* a log-ratio would be the one quantity in the book that did not compose with the others, and the four-register convergence (CANON_INTERNAL §1) would break at exactly the point where feedback enters. **The log-ratio is not a stylistic choice about the**

echo; it is the only form that keeps the echo in the same register as everything the echo has to talk to.

43.0.3 Why the ratio and not the difference

There are two candidate difference measures and the choice between them is decidable, not aesthetic.

```
additive     e_t = M'_t − M_t
ratiometric  rho_t = ln(M'_t / M_t)
```

Three reasons the ratiometric form wins, in increasing order of force.

One: composition. Echoes chain. You send a moment, the world answers, you send the answer onward, a second answer returns. The total return over a chain of k links is

```
M'_k / M_0 = (M'_1/M_0) ∗ (M'_2/M'_1) ∗ ... ∗ (M'_k/M'_{k−1})

ln(M'_k / M_0) = rho_1 + rho_2 + ... + rho_k
```

Ratios compose by multiplication, so log-ratios compose by addition, and a chained return difference is therefore a plain sum. The additive form has no such property: `M' − M` over a chain depends on the absolute magnitudes at every intermediate link, and two chains with the same per-link proportional effect give different totals if they start at different sizes. A measure whose composition law depends on where you started is not measuring the exchange.

Two: scale invariance. Multiply every magnitude in an echo log by a constant — change the modulus of the magnitude-estimation procedure, rate the week out of 100 instead of out of 50 — and every `rho` is unchanged while every additive `e` is scaled. Since the modulus in magnitude estimation is arbitrary by construction, **a difference measure that moves when the modulus moves is measuring the procedure, not the world.**

Three: the lens. CANON_INTERNAL §2, the Lens Theorem: perception is logarithmic, so equal perceptual steps are equal world ratios. The operator reading the echo is a logarithmic instrument. `rho` is what that instrument natively reports. The additive form would require the operator to perform a conversion their perceptual system does not perform, and every such conversion is a place where a bias enters.

43.0.4 Book One, Node 08 ECHO — stated exactly

Book One's eighth node taught the operator to read the difference: that the answer coming back is never the voice sent out, that the gap between them is the only new thing in the exchange, and that two specific pathologies wait on either side of the reading — the Mirror that Flatters, which returns your own voice and calls it agreement, and the Echo that Devours, which returns your voice amplified until it is the only thing in the room.

In this manual, Node 08 is `rho_t = ln(M'_t/M_t)`, its deadband, and its stability condition. Each of the node's three teachings has an exact counterpart:

Book One's teaching	This chapter's quantity	Section
The answer is never the voice sent	`rho` is generically nonzero; `rho = 0` is a measure-zero coincidence	§43.1
The difference is the only data	The magnitude `Delta_R`, its deadband `ln(1+c)`, and its signal-to-noise	§43.2–43.4
The Mirror that Flatters	`rho_obs = (1 − w) rho_true`; correlated echoes and the design effect	§43.5
The Echo that Devours	`rho_{t+1} = (1 − lambda) rho_t` with `|1 − lambda| >= 1`	§43.6

The last row is the one that repays the translation most, because Book One could warn about the devouring echo but could not say where it starts. This chapter can: **it starts at a correction gain of exactly 2, and the marginal case at gain 2 is an undamped oscillation that never grows and never settles.** That is a number, it is derived below, and it is checkable against your own records.

43.0.5 The zero clause, stated hard

CANON_INTERNAL §8 says it in six words: *zero means no information came back.* The manual states the consequence at full strength because it is the most commonly ignored reading in the whole formulary.

`Delta_R = 0` is not success. It is the absence of an exchange. A world that returns exactly what you sent has told you nothing about itself — it has told you only that whatever it did to your voice, it did symmetrically enough to cancel. There are two ways to produce a persistent zero and they are opposite in kind: perfect calibration, in which the operator's construction of the moment already matched what the world would return, and perfect closure, in which the channel back is not the world at all but a reflection of the operator. **The arithmetic cannot tell these apart from a single reading, and §43.5 gives the test that can.**

The corollary is the one that costs people years: **a world that only ever agrees with you is telling you nothing, and the pleasantness of the sensation is not evidence against this.** The sensation of being agreed with and the sensation of being correct are, from the inside, the same sensation. This manual has said the equivalent thing at group scope already — Chapter 11 on pseudodiversity, Chapter 18 on collective intelligence — and Volume Two's thesis (CANON_INTERNAL §0) is that it is one statement at two scopes, not two statements.

43.0.6 Formula Validation Protocol — the return difference

Per CANON_INTERNAL §6, no formula appears in Volume Two without all seven gates in print.

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M` and `M'` carry the units of the magnitude-estimation modulus (arbitrary but fixed). `M'/M` is dimensionless. `ln` of a dimensionless number is dimensionless; `rho` and `Delta_R` are reported in **nats**, convertible to bits, octaves and cents by the CANON_INTERNAL §1 keys. Both sides agree.
FVP-2 SCALE TYPE	**CONDITIONAL PASS — the binding gate**	Division requires ratio scale with a true zero. Magnitude estimation supplies this (Stevens 1957); a 1–5 satisfaction rating does not. **If `M` and `M'` are ordinal, `ln(M'/M)` is not defined and the formula must be demoted in print to a heuristic ordering: "the return was larger / the same / smaller."** The demotion is not optional and §43.8 specifies it.
FVP-3 LIMITING CASES	PASS	`M' -> M`: `rho -> 0`, no information returned (§43.0.5). `M' >> M`: `rho -> +inf`, over-return; in practice bounded by the modulus ceiling. `M' -> 0`: `rho -> -inf`, total non-return — the answer that constructs no moment at all, which is the real phenomenon of being ignored, and the divergence is correct because there is no finite penalty scale for it. `M -> 0`: undefined; correctly so, because nothing was sent and an echo of nothing is not an echo.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	`rho = 0`, i.e. `M' = M`. Under the correction dynamics of §43.6 this is the unique fixed point of the map `rho -> (1-lambda)rho` for any `lambda != 0`, and it is attracting iff `0 < lambda < 2`. Meaning: **the state in which your construction of the moment and the world's answer to it agree — which is either contact or closure, and only the correlation test distinguishes them.**
FVP-5 IDENTIFIABILITY	**PARTIAL**	`rho` is recoverable from paired magnitude estimates. The *decomposition* of `rho` into "the world's contribution" and "the operator's mis-framing" is **not identified from `rho` alone** — a large negative `rho` is equally consistent with an over-framed operator (`r = D/P > 1`, Chapter 35) and with a genuinely unresponsive world. Separating them requires a second instrument: either an independent estimate of `M'` by someone who did not send `M` (§43.5) or a manipulation of `D` with `E` held fixed (§43.10). **State this limitation whenever a single `rho` is used to draw a conclusion about a person.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Paper E §5.3 template. Preregister: the operationalization of `M` and `M'`, the modulus, the correction gain `lambda`, the horizon. Fit `lambda` on a training block, **freeze**, and predict held-out `rho`. Four baselines: (i) persistence — next `rho` equals last `rho`; (ii) direct-input — `rho` predicted from `M` alone; (iii) matched autoregressive AR(1) on `rho`; (iv) **shuffled-input control** — the same `M` series paired with a permuted `M'` series. **The formula dies if the frozen `lambda` model fails to beat the shuffled-input control on held-out data**, because that is the outcome in which the pairing between what was sent and what came back carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	The cross-individual gate. Same operationalization, different operators, different investigators: do the fitted `lambda` values recover, and do they rank people consistently? This manual has not run that study and does not have the data. **What would need measuring is stated in §43.11 as a protocol, and until it is run, `lambda` is a within-person descriptive parameter and nothing more.** This is the gate CANON_INTERNAL §6 says is most often skipped; it is not skipped here, it is declared open.

The Engineering

43.1 Reading one echo: sign first, magnitude second

A single echo yields two readings and they answer different questions.

```
    rho > 0   OVER-RETURN    the world answered larger than you sent
    rho = 0   NULL RETURN    no information came back (contact or closure --
test required)
    rho < 0   UNDER-RETURN   the world answered smaller than you sent
```

The sign says which way you were wrong. The magnitude says how much it cost to find out.

Over-return. You sent a moment and got back more than you constructed. The honest readings, in order of frequency: your frame under-called the environment (`r = D/P < 1`, under-framed in the sense of Chapter 35); or the environment contained structure you had not modelled; or the answer was amplified in transit by something with its own energy — an audience, an institution, a grievance. **Over-return is the pleasant failure and it is still a failure of calibration.** The specific danger is that over-return is usually attributed to the quality of what was sent, which is the one explanation the measurement cannot support.

Under-return. You sent a moment and the world answered smaller. Readings: your frame over-called the environment (`r > 1`, over-framed — projection, expectation, the seen-what-was-predicted failure); or the channel attenuated; or the environment genuinely had less in it than you built into your construction of it. **Under-return is the common case for anyone with a strong frame, and its characteristic error is to be read as a verdict on worth rather than as a reading on calibration.** §43.8 treats this at length because it is where the instrument does the most harm when misused.

Null return. Treated in §43.0.5 and tested in §43.5.

43.2 The magnitude scale — what a given `Delta_R` is

`rho` is in nats. Because the four registers converge (CANON_INTERNAL §1), the same number is immediately readable as a ratio, as bits, and as an interval in octaves and cents. The conversion constants are exact:

```
1 nat = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
```

`Delta_R` (nats)	ratio `M'/M`	percent change	bits	octaves	cents
0.0050	1.0050	0.50%	0.0072	0.0144	17.3
0.0100	1.0101	1.01%	0.0144	0.0289	34.6
0.0200	1.0202	2.02%	0.0289	0.0577	69.2
0.0500	1.0513	5.13%	0.0721	0.1443	173.1
0.1000	1.1052	10.52%	0.1443	0.2885	346.2
0.2000	1.2214	22.14%	0.2885	0.5771	692.5
0.2500	1.2840	28.40%	0.3607	0.7213	865.6
0.3466	1.4143	41.43%	0.5000	1.0001	1200.1
0.5000	1.6487	64.87%	0.7213	1.4427	1731.2

Delta_R (nats)	ratio M'/M	percent change	bits	octaves	cents
0.6931	2.0000	99.99%	0.9999	2.0000	2399.8
1.0000	2.7183	171.83%	1.4427	2.8854	3462.5
1.3863	4.0000	300.00%	2.0000	4.0000	4800.0
2.0000	7.3891	638.91%	2.8854	5.7708	6924.9
2.3026	10.0001	900.01%	3.3219	6.6439	7972.7

Two landmarks are worth memorising because they anchor every subsequent judgement.

Delta_R = 0.3466 nats is exactly one octave and exactly half a bit — the world answered about 1.41 times what you sent, or about 0.71 times. **Delta_R = 0.6931 nats is exactly one bit and exactly two octaves** — a factor of two, either way. A return difference of one bit means the exchange resolved one binary question you did not previously have answered, and if that sounds like a small yield for an entire cycle of action, it is: **exchanges are expensive and their information yield is measured in single bits, which is the correct order of magnitude and the reason patience is an engineering requirement rather than a virtue.**

A caution that must travel with this table, because the table invites the error. **The nats of a return difference are not automatically nats of information.** They become information only relative to the echo's own noise, and §43.4 gives the ratio that decides it. A Delta_R of 0.20 against a measurement noise of 0.30 is not one-fifth of a bit of anything; it is a number smaller than the wobble of the instrument that produced it.

43.3 The deadband — the smallest echo you can hear

Weber's law (1834): the just-noticeable difference is a constant ratio, dI/I = c. Applied to the echo, this gives a hard floor on Delta_R below which the operator's own instrument cannot resolve the return.

```
Delta_R_min = ln(1 + c)
```

The derivation is one line: the smallest ratio the operator can distinguish from unity is 1 + c, and Delta_R is the log of the ratio.

Weber fraction c	Delta_R_min (nats)	bits	ratio	cents
0.01	0.0100	0.0144	1.01	34.5
0.02	0.0198	0.0286	1.02	68.6
0.03	0.0296	0.0426	1.03	102.3
0.05	0.0488	0.0704	1.05	168.9
0.08	0.0770	0.1110	1.08	266.5
0.10	0.0953	0.1375	1.10	330.0
0.15	0.1398	0.2016	1.15	483.9
0.20	0.1823	0.2630	1.20	631.3
0.30	0.2624	0.3785	1.30	908.4
0.50	0.4055	0.5850	1.50	1403.9

The engineering consequence: **an operator with a Weber fraction of 0.20 on moment magnitude cannot distinguish a genuine 15% return from silence, and will experience nine weeks of real but small feedback as nine weeks of nothing.** This is not a failure of attention. It is the resolution of the instrument, and the fixes are instrumental rather than motivational: increase the modulus (rate on a finer scale), aggregate (§43.4), or amplify the sent moment so the return clears the floor.

Note also the honest reading of the deadband in the other direction. **A null return inside the deadband is not evidence of a null return.** It is evidence that the echo, whatever it was, was smaller than you can hear. The distinction matters because the two conclusions license opposite actions: *no signal* says change what you are sending; *below resolution* says change how you are listening.

An empirical caveat with its badge: this manual has not measured a Weber fraction for moment magnitude and does not know one. **What would need measuring is stated in §43.11.** The table above is the arithmetic of the relation, not a claim about any particular person's threshold.

43.4 Noise, aggregation, and how many echoes a finding costs

A single `rho` is a noisy quantity. Estimate the noise as `sigma_rho`, the standard deviation of repeated log-ratios under conditions you regard as unchanged. Then the echo's signal-to-noise is

```
SNR_R = Delta_R / sigma_rho
```

Delta_R \ sigma_rho	0.05	0.10	0.20	0.30	0.50
0.05	1.00	0.50	0.25	0.17	0.10
0.10	2.00	1.00	0.50	0.33	0.20
0.20	4.00	2.00	1.00	0.67	0.40
0.35	7.00	3.50	1.75	1.17	0.70
0.50	10.00	5.00	2.50	1.67	1.00
0.69	13.80	6.90	3.45	2.30	1.38
1.00	20.00	10.00	5.00	3.33	2.00

Aggregation reduces noise in the usual way — the standard error of a mean of `n` independent log-ratios is `sigma_rho / sqrt(n)` — so the number of echoes needed to resolve a true return difference `Delta` at the conventional two-sided 95% level is

```
n >= ( 1.96 * sigma_rho / Delta )^2
```

sigma_rho \ Delta	0.05	0.10	0.20	0.35	0.50	0.69
0.10	16	4	1	1	1	1
0.20	62	16	4	2	1	1
0.30	139	35	9	3	2	1
0.50	385	97	25	8	4	3

Read the top-left corner and the bottom-left corner together: a small return difference under large noise is not a finding you can afford, and the honest response is to stop trying to detect

it rather than to detect it badly. An operator with `sigma_rho = 0.30` who wants to establish a 5% return effect needs 139 clean, independent echoes. If the cycle time is a week, that is two and a half years. The engineering answer is never "look harder"; it is to increase `Delta` (send a moment big enough to produce a hearable answer), reduce `sigma` (fix the estimation procedure), or abandon the question.

The word *independent* in that paragraph is doing a great deal of work, and the next section is about what happens when it is false.

43.5 The Mirror that Flatters — correlated echoes and the design effect

Book One's first echo pathology, in exact form.

Suppose a fraction `w` of the answer that comes back is not the world's response but a reflection of the moment you sent — the audience that has learned your frame, the friend who reports what you want, the metric you designed. Model it linearly: the observed return is a blend of the true return and your own sent moment. Then

```
rho_obs = (1 − w) ∗ rho_true
```

w (reflected fraction)	attenuation	rho_true required to observe rho_obs = 0.10
0.00	1.00	0.1000
0.10	0.90	0.1111
0.20	0.80	0.1250
0.30	0.70	0.1429
0.50	0.50	0.2000
0.70	0.30	0.3333
0.90	0.10	1.0000
0.95	0.05	2.0000
1.00	0.00	infinite — no echo exists

At `w = 1` the channel is a mirror and no finite world could produce a detectable return: the apparatus reports agreement no matter what is true. That is the flattering mirror, and it is not a metaphor in this form — it is a divide-by-zero in the recovery of `rho_true`.

The multi-source version is more useful because most people have several echo sources and believe that several is the same as independent. It is not. Where `n` echo sources share an intraclass correlation `r_icc`, the effective number of independent echoes is the standard design effect:

```
n_eff = n / ( 1 + (n − 1) ∗ r_icc )
```

For `n = 10` sources:

r_icc	n_eff	standard-error multiplier vs. 10 independent	Paper A's N_eff = 1 + (n−1)(1−r) for comparison
0.00	10.000	1.000	10.000
0.05	6.897	1.204	9.550

r_icc	n_eff	standard-error multiplier vs. 10 independent	Paper A's N_eff = 1 + (n−1)(1−r) for comparison
0.10	5.263	1.378	9.100
0.20	3.571	1.673	8.200
0.30	2.703	1.924	7.300
0.50	1.818	2.345	5.500
0.70	1.370	2.702	3.700
0.90	1.099	3.017	1.900
1.00	1.000	3.162	1.000

Ten sources at an intraclass correlation of 0.30 are worth 2.7 independent sources, and the person consulting them experiences ten confirmations. That gap — between the felt count and the audited count — is the same gap Paper A's effective count measures for group members, and the last column is included to make an honest point rather than a tidy one: **the two audits agree exactly at both endpoints and disagree in between.** `N_eff = 1 + (n−1)u` and `n_eff = n/(1+(n−1)r_icc)` both collapse to 1 under total redundancy and both return `n` under total independence, but they are different functions of the correlation and they are answering different questions — Paper A prices the surplus generated by a set of states, the design effect prices the precision of a mean. Do not substitute one for the other. **Use the design effect when you are averaging echoes and Paper A when you are asking what the echoes are worth as a set.**

The zero-return test. This is the instrument §43.0.5 promised — the way to tell contact from closure when `Delta_R` is persistently near zero. Under genuine contact with a live world, observed `rho` values scatter: the return is sometimes over, sometimes under, by amounts comparable to the measurement noise. Under closure, they do not scatter, because the channel is returning the sent moment. So compute the probability that `n` genuinely independent echoes all land inside a deadband of half-width `b` when the true return difference is zero and the noise is `sigma`:

sigma	band b	P(one echo inside band)	P(10 in a row)	n for p^n < 0.05
0.10	0.05	0.3829	0.000068	4
0.10	0.10	0.6827	0.021990	8
0.20	0.05	0.1974	0.0000000	2
0.20	0.10	0.3829	0.000068	4
0.30	0.05	0.1324	0.0000000	2
0.30	0.10	0.2611	0.0000010	3

Read the right-hand column as the flattery alarm: with a noise level of 0.20 and a deadband of 0.10, four consecutive echoes landing inside the band is already a 5%-level event under the hypothesis that the channel is live and merely quiet. Four weeks of "that was great" is not warmth. It is a measurement, and the measurement says the channel is closed.

Two honest limits on this test. First, it assumes the noise `sigma` was estimated on the same channel under conditions where the channel was known to be live, which requires a period of record before the suspicion arose. Second, it cannot distinguish closure from genuine perfect

calibration, and perfect calibration does occur — a practitioner who has converged genuinely does get small, scattered returns. **The discriminator is scatter, not size: converged contact produces small returns that vary; closure produces small returns that do not.**

43.6 The Echo that Devours — correction gain and the stability boundary

Book One's second echo pathology, and the place where this chapter produces its sharpest number.

An echo is only useful if it changes something. The natural thing for it to change is the frame: the operator's perspective `D` in the Moment equation `M = (D/P)E` (Chapter 34). Correct multiplicatively, because the register is multiplicative:

```
D_{t+1} = D_t * ( M'_t / M_t )^lambda        lambda >= 0, the CORRECTION GAIN
```

Now trace what that does to the next return difference. Write `A` for the answer the world would give — the quantity the operator is trying to match — and let the sent moment be `M_t = (D_t / P) E`. Then `rho_t = ln(A / M_t)`, and

```
M_{t+1} = (D_{t+1}/P) E = M_t * (A/M_t)^lambda

rho_{t+1} = ln(A / M_{t+1})
          = ln(A/M_t) - lambda * ln(A/M_t)
          = (1 - lambda) * rho_t
```

The return difference obeys a linear contraction in the log register with factor `(1 - lambda)`, and that is the same mathematical object as the storage law's `(1 - delta)` — which is why this chapter is the setup for the next one. The echo and the accumulator are the same circuit read at two places. Chapter 44 works the consequence out in full; here is the stability table.

lambda	|1 − lambda|	behaviour	steps to reduce |rho| 10x	steps to reduce 100x
0.10	0.9000	monotone, slow	22	44
0.20	0.8000	monotone	11	21
0.25	0.7500	monotone	9	17
0.33	0.6700	monotone	6	12
0.50	0.5000	monotone	4	7
0.75	0.2500	monotone, fast	2	4
1.00	0.0000	**deadbeat** — one step	1	1
1.25	0.2500	oscillatory, converging	2	4
1.50	0.5000	oscillatory	4	7
1.75	0.7500	oscillatory, slow	9	17
2.00	1.0000	**marginal** — undamped oscillation	never	never
2.50	1.5000	**divergent** — the devouring echo	never	never

Five statements follow, and they are the operational content of Node 08's second warning.

One. There is a stability boundary and it is at `lambda = 2`. Below it the return difference shrinks; above it the return difference grows without bound. At exactly 2 it oscillates forever at constant amplitude. This is the Echo that Devours, located precisely: it is not an atmosphere, it is a gain setting.

Two. `lambda = 1` is the fastest possible correction and it is a trap. Deadbeat correction takes the whole observed difference as truth and applies it in one step. If the observation were noiseless this would be optimal. It is not noiseless, so `lambda = 1` also takes the whole observation *noise* as truth and writes it into the frame. The frame then thrashes at the amplitude of the measurement error.

Three. Everything above 1 is oscillatory even when it converges. An operator at `lambda = 1.5` overshoots every time and converges by alternating: too much, too little, too much. From the inside this is indistinguishable from instability of character, and it is a gain setting.

Four. Slow gains are not timid; they are filters. `lambda = 0.25` needs nine cycles to close a factor of ten, and in exchange it attenuates the noise in every single observation by 75%. The choice of `lambda` is a bandwidth choice: **fast correction tracks the world and the noise; slow correction tracks the world only, late.** There is no setting that does both, and the belief that there is one is the most common source of thrash.

Five. The optimum depends on the noise, not on the character of the operator. As a design rule, badged as a heuristic rather than a derived optimum: set `lambda` near `SNR^2/(1+SNR^2)` where `SNR` is the per-echo signal-to-noise from §43.4 — near 1 when echoes are clean, near 0 when they are not. This manual has not derived that expression from a stated loss function and does not present it as optimal; it presents it as a starting value with the right limits.

43.7 Delay, aliasing, and the echo that arrives after the world has moved

Returns take time. If the world's answer to period `t` arrives at period `t + k`, and the operator applies it as if it were the answer to period `t + k`, the correction is applied to the wrong state. In control terms this is transport delay, and its effect is to convert a stable gain into an unstable one: a correction that would be safely damped at zero delay can oscillate or diverge at a delay of several periods, because the operator keeps correcting for a discrepancy that has already been corrected.

The practical rule, and the reason it appears in a manual about persons rather than a manual about servos: **never run a correction gain faster than the echo returns.** If the answer takes three weeks, do not correct weekly. Correct every three weeks, or reduce `lambda` by roughly the delay factor, or — best — timestamp the echo to the moment it answers rather than to the moment it arrives. The third option costs nothing but a column in the log and removes the problem at its source, which is why the specification below requires it.

The dual failure is sampling too slowly. If the world's answer changes faster than you sample it, the sampled series aliases: a rapid alternation reads as a slow trend or as nothing at all. The symptom is a return log whose variance is implausibly low for a domain everyone agrees is volatile. **A suspiciously smooth echo log is a sampling artifact until proven otherwise.**

43.8 What the echo is not

Four things, because the instrument is dangerous in proportion to how naturally it is misread.

It is not a verdict on worth. `rho < 0` says the world answered smaller than you constructed. It does not say the thing was bad, and FVP-5 above is the formal statement of why: a large negative `rho` is equally consistent with an over-framed operator and an unresponsive world, and `rho` alone cannot separate them. Anyone who reads a run of negative returns as a statement about their worth has performed an identification that the mathematics explicitly declines to perform.

It is not defined on ordinal data. FVP-2 is a gate, not a caution. If `M` and `M'` are 1–5 stars, `M'/M` is meaningless — the ratio of two rank positions is not a ratio of two magnitudes. The demotion in that case is mandatory and it is this: report the sign only, as an ordering (*larger* / *same* / *smaller*), call it a heuristic reading, and **do not compute a mean, a standard deviation, a `lambda`, or a convergence rate from it.** A great deal of self-tracking practice violates this gate silently, and the resulting numbers have the shape of measurement without the content.

It is not the world's total answer. As deployed, `rho` is one scalar summary of an exchange that had structure. This mirrors CANON §1.1's honesty clause on the surplus functional — we log the floor, not the ceiling. **A `rho` of zero on the dimension you measured is compatible with an enormous return on a dimension you did not.**

It is not transferable between operators without recalibration. `M` and `M'` are magnitude estimates made against a modulus. Two people using different moduli produce `rho` values that are comparable — this is the scale-invariance property of §43.0.3 — but only if each person's own estimates are internally consistent on a ratio scale. The comparison of `lambda` across people is the open FVP-7 gate, not a licensed operation.

43.9 Chaining and the composition of returns

Because log-ratios add, a series of exchanges has a total return difference that is the sum of the signed parts and a total *activity* that is the sum of the magnitudes. These are different numbers and their ratio is a diagnostic.

```
net return   R_net = SUM rho_t                        (signs cancel)
gross return R_gross = SUM |rho_t| = SUM Delta_R,t   (signs do not)

EFFICIENCY   E_R = |R_net| / R_gross      in [0, 1]
```

`E_R` near 1 means the returns pointed consistently the same way: the world has been saying one thing repeatedly and the operator has not yet moved. `E_R` near 0 means the returns cancelled: a great deal of exchange, no net direction. **A high gross with a near-zero net is the signature of thrash — the operator is correcting hard in alternating directions, which §43.6 identifies as a gain above 1.** A high gross with a high net is the signature of a live, uncorrected error, which is the good case, because it means there is something to learn and it has not been learned yet.

43.10 The join: from echo to stored order

The formulary is a single circuit and this section names the wires. The chain runs:

```
E (environment)
   |
```

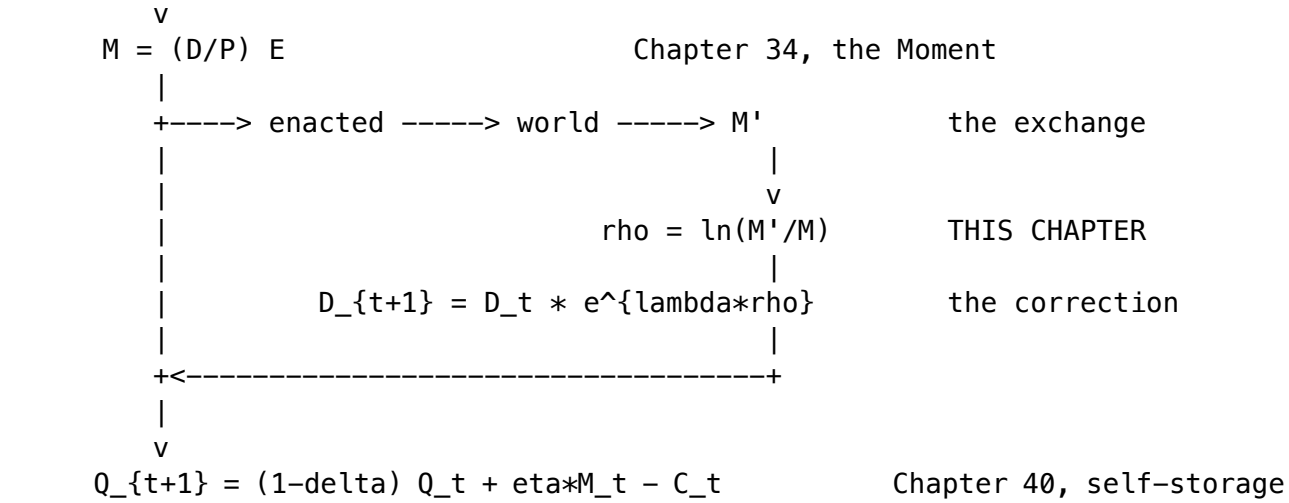

Three consequences of drawing it this way.

The echo does not enter the storage law directly. It enters through `D`, which changes `M`, which is the storage law's input. This is why feedback improves character slowly: it acts on the input of an integrator, not on the integrator's state. An echo received today changes `Q`'s trajectory, not `Q`.

The correction loop and the storage loop have different time constants and both are geometric. The echo loop contracts at `(1 − lambda)`, the storage loop at `(1 − delta)`. When `lambda >> delta` the frame tracks and the character lags; when `lambda << delta` the character forgets faster than the frame corrects, which is the arithmetic form of learning a lesson repeatedly.

The identification problem of FVP-5 has a design solution. To separate "my frame was wrong" from "the world was quiet", hold `D` fixed for a block of periods and vary nothing, then change `D` deliberately by a known factor and see whether `rho` moves by the predicted amount. If `rho` responds to a change in `D` with the elasticity the Moment equation predicts, the return difference is reading your frame. If it does not, it is reading the world. **That is a manipulation, it is cheap, and it is the only clean way to run the decomposition.**

43.11 What would need measuring

CANON §2 requires that where a figure needs research not done, the manual says what would need measuring. Four things, stated as an agenda rather than as results:

1. **The Weber fraction for moment magnitude.** Method of constant stimuli or a staircase procedure on paired moment magnitudes within a person, yielding `c` and hence the deadband `ln(1+c)`. Without it, §43.3 is arithmetic without a calibration.
2. **The within-person noise `sigma_rho`.** Test–retest of `rho` on exchanges the operator judges equivalent. Everything in §43.4 depends on it and it is the cheapest of the four to obtain.
3. **The distribution of `lambda` across people, and its stability within a person across domains.** This is FVP-7. The design is stated: same operationalization, several operators, several investigators, held-out prediction, shuffled-input control.
4. **The intraclass correlation of a typical person's echo sources.** This determines whether the flattering-mirror alarm in §43.5 fires at four echoes or at forty, and it is measurable by having several sources rate the same sent moment independently.

The Specification

Numbered, checkable requirements for an echo instrument. A practice that fails any of ES-1 through ES-6 is not measuring return difference and should not report a number.

ES-1 — Ratio scale, or demotion. `M` and `M'` shall be obtained by magnitude estimation against a stated modulus, with a true zero. If only ordinal data are available, the practice shall report sign only and shall label the output a heuristic ordering. *(FVP-2. Non-negotiable.)*

ES-2 — The modulus is declared and frozen. One reference exchange is assigned a magnitude (commonly 50 or 100) at the start of a block, and all estimates in the block are made against it. Changing the modulus mid-block invalidates the block for aggregation.

ES-3 — `M` is recorded before the answer arrives. The sent moment shall be estimated and written down prior to any return. An `M` reconstructed after the echo is contaminated by the echo and yields a `rho` biased toward zero. *(This is the single most violated requirement in self-tracking and it is the one that most reliably manufactures the appearance of calibration.)*

ES-4 — `M'` is timestamped to the exchange it answers, not to its arrival. A delay column is mandatory. *(§43.7.)*

ES-5 — Source and independence are logged. Each echo carries the identity of its source and, where more than one source is used, sources are asked independently and before seeing one another's responses. *(§43.5; this is the personal-scope form of Chapter 17's silent-write-first rule.)*

ES-6 — Noise is estimated before findings are claimed. `sigma_rho` is estimated from at least eight exchanges the operator judges equivalent, before any `Delta_R` is interpreted as a finding.

ES-7 — The deadband is stated. The practice declares its `Delta_R_min` and refuses to interpret returns below it as anything but unresolved.

ES-8 — The correction gain is chosen, written down, and held for the block. `lambda` shall be in `(0, 1]` unless a specific reason for overshoot is recorded. `lambda >= 2` is prohibited by this specification. *(§43.6.)*

ES-9 — Correction cadence does not exceed echo cadence. If returns take `k` periods, the gain is applied at most every `k` periods. *(§43.7.)*

ES-10 — The flattery alarm is run. After every four consecutive returns inside the deadband, the practice runs the zero-return test of §43.5 and, if it fires, adds an independent source before continuing.

ES-11 — Net and gross are both reported. `R_net`, `R_gross`, and `E_R` are computed per block. *(§43.9.)*

ES-12 — No worth claims. The instrument's output is a calibration reading. Reports shall not convert `rho` into a statement about the value of the person, the work, or the relationship. *(§43.8, and CANON §5.)*

ES-13 — Decline path. Where an echo source is a person, that person may decline to supply an estimate, at any time, without explanation and without consequence, and the practice records the decline as missing data rather than as a zero. *(CANON §5.4 is binding: no protocol in this*

manual removes a person's ability to decline. A refused echo is not a null return; recording it as one is both a measurement error and a coercion.)

ES-14 — Freeze before predicting. Any claim that the instrument works shall follow the FVP-6 procedure: preregister, fit, freeze, predict held out, beat the four baselines including the shuffled-input control.

The Failure Modes

Fault	Diagnostic signature	Fix
The Flattering Mirror	Four or more consecutive returns inside the deadband, with variance far below `sigma_rho`; all sources agree; `E_R` undefined because gross is near zero	Add a genuinely independent source (ES-5, ES-10). Compute `n_eff`. Treat the existing sources as one source until proven otherwise
The Devouring Echo	`|rho|` growing period on period; corrections alternating in sign with increasing amplitude; frame reversals faster than the world could plausibly change	Measure `lambda` from the data (regress `ln|rho|` on `t`). If the implied `|1−lambda| >= 1`, reduce `lambda` below 1 and hold. *(§43.6)*
Ordinal laundering	Means, SDs and trend lines computed from star ratings, tick-boxes, or Likert items	Demote to sign-only reporting; rebuild the instrument on magnitude estimation. **The numbers already computed are void, not merely imprecise** *(FVP-2, ES-1)*
Post-hoc `M`	`rho` unusually close to zero across the whole log; the operator "knew it would land like that"	Enforce ES-3. Compare a block of pre-recorded `M` against a block of reconstructed `M`; the difference in mean `|rho|` is the size of the contamination
Deadband blindness	Long runs of "no feedback" in a domain where others report feedback	Raise the modulus resolution, aggregate per §43.4, or increase the size of the sent moment. Do not increase effort at listening
Delay-induced instability	Oscillation that appears only when cadence is increased; corrections that were stable at monthly review become unstable at weekly	Enforce ES-4 and ES-9. Timestamp to the exchange, not the arrival

Fault	Diagnostic signature	Fix
Aliasing	Return log with implausibly low variance in a volatile domain; long smooth trends with no reversals	Sample faster than the world changes, or state explicitly that the log is a low-pass view and stop reading its trend
Selection of echoes	The log contains more returns from the exchanges the operator felt good about; missing entries cluster	Pre-commit to logging every exchange in a defined class, including the ones you would rather not; count and report the missing
Worth conversion	Negative `rho` narrated as a statement about the self; the log stops being kept after a bad block	Enforce ES-12. Re-read FVP-5: the decomposition is not identified, and the identification being performed is not licensed by the instrument
Zero read as success	`Delta_R ~ 0` celebrated; the practice concludes it has arrived	Run the scatter test (§43.5). Converged contact produces small returns that **vary**; closure produces small returns that do not
Gain thrash	High `R_gross`, near-zero `R_net`, `E_R < 0.2`, frequent reversals	`lambda` is above 1. Reduce it. *(§43.9, §43.6)*
Uncorrected live error	High `R_gross`, high `R_net`, `E_R > 0.8`, all returns pointing the same way for many periods	`lambda` is near 0 — the operator is measuring and not correcting. Raise `lambda` toward 0.3–0.5

Worked Example

A twelve-week echo log. The data below are **illustrative, not empirical**: they are a constructed log used to demonstrate the computations, and no claim is made that they describe any real person. Every number derived from them was computed, not estimated by eye.

Setting. An operator produces one unit of work per week in a domain where a return arrives within the same week. Before releasing each week's work, they record `M` — a magnitude estimate of the moment they believe they are sending, on a scale where a fixed reference exchange from the preceding quarter was assigned 50. After the return arrives, they record `M'` on the same scale, from three sources asked independently and before seeing one another (ES-5), averaged. Their declared deadband is `Delta_R_min = 0.05` and their pre-estimated noise from an eight-exchange calibration block is `sigma_rho = 0.09`.

The log

Week	M sent	M' returned	ratio	rho	Delta_R	cumulative mean rho
1	40	22	0.5500	-0.5978	0.5978	-0.5978
2	45	26	0.5778	-0.5486	0.5486	-0.5732
3	50	34	0.6800	-0.3857	0.3857	-0.5107
4	48	35	0.7292	-0.3159	0.3159	-0.4620
5	55	44	0.8000	-0.2231	0.2231	-0.4142
6	52	45	0.8654	-0.1446	0.1446	-0.3693
7	60	55	0.9167	-0.0870	0.0870	-0.3290
8	58	57	0.9828	-0.0174	0.0174	-0.2900
9	55	56	1.0182	+0.0180	0.0180	-0.2558
10	60	63	1.0500	+0.0488	0.0488	-0.2253
11	62	60	0.9677	-0.0328	0.0328	-0.2078
12	60	61	1.0167	+0.0165	0.0165	-0.1891

Block statistics

```
all 12 weeks     mean rho = −0.1891    sd = 0.2257
weeks 1–6        mean rho = −0.3693    sd = 0.1785
weeks 7–12       mean rho = −0.0090    sd = 0.0478

mean Delta_R, weeks 1–6   = 0.3693
mean Delta_R, weeks 7–12  = 0.0368
```

With six observations per half-block, the standard errors and 95% intervals are:

```
weeks 1–6    SE = 0.0729    95% CI on mean rho = [−0.5121, −0.2264]
weeks 7–12   SE = 0.0195    95% CI on mean rho = [−0.0473, +0.0293]
```

The first half-block's interval excludes zero decisively; the second half-block's interval contains zero. The instrument has detected a real, large, consistently signed return difference in weeks 1–6 and cannot detect one in weeks 7–12.

Reading the sign

Every return in weeks 1–6 is negative and none is inside the deadband. This is sustained **underreturn**: the world answered systematically smaller than the operator constructed. Per §43.1, the leading candidate is over-framing — `r = D/P > 1`. The size of the implied over-frame is

```
implied over-frame factor = exp( −mean rho over weeks 1–6 )
                          = exp( 0.3693 )
                          = 1.4467
```

The operator was constructing the moment about 1.45 times larger than the world was returning it. That is a specific, checkable, correctable claim, and it is the sort of claim the additive difference measure could not have produced — `M' − M` in weeks 1–6 runs -18, -19, -16, -13, -11, -7, which looks like a shrinking problem, while the ratio shows a shrinking *frame error*, which is a different fact.

Estimating the correction gain from the data

The operator was correcting — not by a written rule, but the log allows the effective gain to be recovered. Per §43.6, if `rho_{t+1} = (1−lambda) rho_t`, then `ln|rho_t|` falls linearly in `t`. Regressing `ln|rho|` on week for weeks 1–8 (the run before the residuals reach the deadband):

```
slope of ln|rho| on t   = −0.4436
per-step factor  g      = exp(−0.4436) = 0.6418
implied gain     lambda = 1 − g        = 0.3582
R^2                     = 0.8587
```

The operator's effective correction gain was 0.36 — comfortably inside the stable, monotone band, which is why the log shows steady shrinkage with no sign reversals until the residuals reach the noise floor. Derived quantities:

```
half-life of the distortion = ln(0.5)/ln(0.6418) = 1.56 weeks
weeks to bring |rho| below the deadband 0.05     = 7
```

The fitted trajectory and the observed one agree to the accuracy the R^2 implies:

```
fitted |rho| at weeks 1..12:
0.913, 0.586, 0.376, 0.241, 0.155, 0.099, 0.064, 0.041, 0.026, 0.017, 0.011, 0.007
```

Observed week 7 is 0.0870 against a fitted 0.0638; observed week 8 is 0.0174 against 0.0409. From week 9 onward the fitted values are below the deadband and the observed values are indistinguishable from noise, exactly as §43.3 predicts. **After week 8 the model is no longer being tested by the data, because the data cannot resolve it.**

Net, gross, efficiency

```
R_net   = sum rho   = −2.2696
R_gross = sum |rho| =  2.4362
E_R     = 0.9316
```

`E_R = 0.93` is very high. Per §43.9 this is the signature of a **live, consistently directed error** rather than thrash — the world was saying one thing, repeatedly, and the operator moved toward it steadily rather than oscillating around it. Compare the counterfactual: had the same gross activity been produced by a gain above 1, the signs would have alternated and `E_R` would have collapsed toward zero.

The flattery check

Weeks 9–12 give four consecutive returns inside the declared deadband of 0.05 — `+0.0180`, `+0.0488`, `−0.0328`, `+0.0165`. ES-10 fires the alarm. Running the test of §43.5 with `sigma = 0.09` and band `b = 0.05`, the probability of one echo landing inside the band under a live-but-null channel is 0.4215, so four in a row has probability 0.0316 — below 5%.

The alarm fires, and it should, and the alarm is a prompt to test rather than a verdict. Here the test is the scatter criterion: the four residuals are `+, +, −, +` with a standard deviation of 0.0289 against a noise estimate of 0.09. **They vary and they change sign.** That is the signature of converged contact, not of a mirror — a mirror returns the sent moment and produces residuals that do not alternate. The operator records the alarm, records the scatter test that resolves it,

and continues. Had the four residuals been `+0.004, +0.006, +0.003, +0.005`, the correct conclusion would have been the opposite one, and the correct action would have been to add a source who had not been asked before.

What the operator does next

Three decisions follow from the block, and each is licensed by a specific number.

1. Stop correcting the frame. Weeks 7–12 have a mean `rho` of -0.0090 with a 95% interval containing zero and a mean `|rho|` of 0.0368, below the declared deadband of 0.05. **There is no resolvable error left to correct at this modulus, and continuing to apply a gain of 0.36 to unresolvable residuals is how a stable operator becomes an oscillating one.** ES-8's instruction is to hold the gain for the block; the block is over, and the next block's gain should be set from the new signal-to-noise, which is now below 1.

2. Increase the sent moment or increase the resolution — but not both blindly. Per §43.4, with `sigma_rho = 0.09` the operator can resolve a `Delta_R` of 0.10 with four echoes and a `Delta_R` of 0.05 with sixteen. If the next question is finer than 0.10, the instrument needs either more echoes per finding or a finer modulus. If it is coarser, the instrument is already adequate and the effort should go elsewhere.

3. Do not conclude anything about the work. ES-12. The block establishes that the operator's construction of the moment moved from about 1.45x the world's return to about 1.0x. It establishes nothing about whether the work is good, and FVP-5 states formally why: the decomposition of `rho` into frame error and world response is not identified from `rho` alone. **The operator has calibrated an instrument. That is the entire finding, and it is worth having.**

The one number to carry out of the chapter

```
weeks 1–6   mean Delta_R = 0.3693 nats = 0.5328 bits = 1.066 octaves
weeks 7–12  mean Delta_R = 0.0368 nats = 0.0531 bits = 0.106 octaves
```

Twelve weeks of exchange bought about half a bit of correction, and half a bit is a real, respectable, hard-won yield. Anyone expecting more from twelve exchanges has mistaken the register they are working in. The world answers in ratios, the operator hears in logarithms, and the difference between the voice and the answer is the only thing either of them learns.

Chapter 44 . The Contraction and Arrival

The Law

Book One's ninth node taught the spiral rather than the circle: repetition that absorbs its own error and returns to the same post a little higher. Its tenth node taught the room that was always furnished — arrival as recognition rather than as victory, a destination that turns out to have been visited many times in imagination before it was reached in fact.

Those are not two teachings. They are the two terms of one equation, and this chapter is where the correspondence is proved rather than asserted.

CANON §1.3 states the correspondence: Book One's Node 10 FIXEDPOINT *is* Paper E, Theorem 3.5, and Node 09 RECURSE *is* the transient `(1−delta)^t Q_0`. Appendix F states it again. Neither place proves it. The claim as it stands is a claim of identity between a ceremonial teaching and a theorem, and a claim of identity has to be discharged, not repeated. **This chapter discharges it: it constructs the metric space, verifies the contraction, invokes Banach, solves for the fixed point, derives the exact error law, and shows that every property Book One attributed to arrival is a property of that solution and not an interpretation laid over it.**

44.0.1 The starting point

Paper E, §3 — the storage law, in its internal-scope form (CANON_INTERNAL §8, the self-storage law):

```
Q_{t+1} = (1 − delta) Q_t + eta ∗ M_t − C_t

Q_t      stored order: character, skill, memory, internalized structure
M_t      the moment as lived, normalized into [0, ln N]
delta    retention loss in [0,1] −− forgetting, atrophy, drift of self
eta      conversion efficiency >= 0 −− how experience becomes character
C_t      cognitive load, effort, maintenance of self, >= 0
```

Take the steady-regime case that Theorem 3.5 addresses: constant geometry `M_t = M∗` for all `t`, constant cost `C_t = c∗`, and `delta > 0`. Define the update map

```
T : R −> R,      T(Q) = (1 − delta) Q + (eta ∗ M∗ − c∗)
```

Write `k = 1 − delta` for the contraction factor and `b = eta∗M∗ − c∗` for the drive. Then `T(Q) = kQ + b`, and the whole of Nodes 09 and 10 is the behaviour of iterating this map.

44.0.2 Step one — the space is complete

Banach's theorem requires a **complete** metric space. Take `X = R` with the usual metric `d(x,y) = |x − y|`. The reals with the absolute-value metric are complete: every Cauchy sequence converges, and its limit is real. This is standard and it is stated here because the proof needs it and because the honest form of the argument names every hypothesis it uses.

One caveat, discharged now. Admissibility (`C_t <= (1−delta)Q_t + eta S_t`) keeps the modelled quantity in `[0, inf)`, which is also complete, so the argument runs on either space. **We run it on `R` deliberately: Theorem 3.5's negative fixed point lives outside `[0, inf)`, and a proof on the restricted space would hide exactly the case the theorem refuses to hide.**

44.0.3 Step two — the map is a contraction, with constant `1 - delta`

For any `Q_1, Q_2 in R`:

```
| T(Q_1) - T(Q_2) |  =  | (kQ_1 + b) - (kQ_2 + b) |
                     =  | k (Q_1 - Q_2) |
                     =  |1 - delta| * | Q_1 - Q_2 |
```

For `delta` in `(0, 1]`, `|1 - delta| = 1 - delta < 1`. The map is Lipschitz with constant `1 - delta`, and that constant is strictly less than 1. **T is a contraction, and it is a contraction for exactly one reason: `delta > 0`.**

Note that the inequality in the definition of a contraction is here an **equality**. T is affine, so it does not merely shrink distances by at most `1-delta`; it shrinks every distance by exactly `1-delta`. This is stronger than Banach requires, and §44.0.6 collects what the extra strength buys.

At `delta = 0` the constant is 1, T is an isometry, and Banach does not apply — correctly, since `T(Q) = Q + b` has no fixed point at all when `b != 0`. That is CANON §1.3's contrast made exact: `Q_t = Q_0 + t*b`, unbounded linear growth, no arrival. **The leak is what makes arrival possible, and the proof shows precisely where: it is the single step above where `1 - delta < 1`.**

44.0.4 Step three — Banach, and the fixed point solved

Banach's fixed point theorem: a contraction on a nonempty complete metric space has exactly one fixed point, and the iterates from any starting point converge to it.

Applied: T has a unique `Q_inf` with `T(Q_inf) = Q_inf`, and `Q_t = T^t(Q_0) -> Q_inf` for every `Q_0`. Solve:

```
Q_inf = (1 - delta) Q_inf + eta*M* - c*
delta * Q_inf = eta*M* - c*

Q_inf = ( eta * M* - c* ) / delta              <-- Paper E, Theorem 3.5
```

which, writing `M* = F(u*) = ln(1 + (N-1)u*)` for the surplus form, is Theorem 3.5 verbatim: `Q_inf = (eta F(u*) - c*) / delta`.

Three things were claimed by Book One's tenth node and all three are now consequences of the line above rather than readings of it.

Existence. Banach guarantees a fixed point. The room is furnished. There is a state the practice converges to, and it exists whether or not anyone has been there.

Uniqueness. Banach guarantees exactly one. There is no second room. Two different practices with the same `(eta, M*, c*, delta)` arrive at the same place, and one practice cannot arrive at two.

Independence of the start. The expression for `Q_inf` **contains no `Q_0`**. Where you arrive does not depend on where you began. Book One said arrival feels like recognition rather than victory; §44.3 shows that this is what the absence of `Q_0` from that formula feels like from inside.

44.0.5 Step four — the exact error law

Banach gives convergence. The affine structure gives the exact rate. Subtract the fixed-point equation from the update:

```
Q_{t+1} − Q_inf = (1−delta) Q_t + b  −  [ (1−delta) Q_inf + b ]
                = (1−delta) ( Q_t − Q_inf )
```

so by induction

```
Q_t − Q_inf = (1 − delta)^t ( Q_0 − Q_inf )

| Q_t − Q_inf | = (1 − delta)^t | Q_0 − Q_inf |        THE ERROR LAW
```

and rearranged, the form that carries the most meaning:

```
Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )
      −−−−−   −−−−−−−−−−−−−−−−−−−−−−−−−−
      where     what remains of where you
      you are   started, decaying at rate delta
      going
```

The entire trajectory is the destination plus a fading memory of the origin. That single line is Nodes 09 and 10 together: Node 10 is the first term, Node 09 is the second, and the reason the manual has been calling them one thing is that they are the two terms of this decomposition.

Consistency check against Paper E, Theorem 3.1 (the exact solution). With constant drive b,

```
Q_t = (1−delta)^t Q_0 + SUM_{j=0}^{t−1} (1−delta)^j ∗ b
    = (1−delta)^t Q_0 + b ∗ (1 − (1−delta)^t)/delta
    = (1−delta)^t Q_0 + Q_inf ∗ (1 − (1−delta)^t)
    = Q_inf + (1−delta)^t (Q_0 − Q_inf)
```

The geometric sum closes exactly. **Theorem 3.1 and the error law are the same statement, and the Banach argument was not needed to obtain it in this one-dimensional affine case — it was needed to know that the result is not an accident of linearity.** That distinction is the subject of §44.0.6.

44.0.6 What Banach buys that the algebra does not

An honest accounting, because it would be easy to over-sell the theorem here.

For this specific affine map on R, the closed form above is elementary and Banach is not required. What Banach supplies is **robustness of the conclusion to changes in the model that the closed form would not survive**:

- If `eta`, `C`, or the moment-generating mechanism are **nonlinear in `Q`** — as they are in any realistic account where effort rises with what must be maintained — the closed form vanishes but the fixed point survives, provided the map remains a contraction. The condition to check becomes `|T'(Q)| <= k < 1` on the region of interest.
- If `Q` is **multi-dimensional** — several stored capacities that interact — the same theorem applies on `R^n` with any complete norm, and the contraction condition becomes a bound on the matrix norm of the update.

- Banach also supplies **error bounds computable before convergence**, which the closed form gives only after you already know `Q_inf`:

```
a priori:       d(Q_t, Q_inf) <= ( k^t / (1-k) ) * d(Q_1, Q_0)
a posteriori:   d(Q_t, Q_inf) <= ( k / (1-k) )  * d(Q_t, Q_{t-1})
```

The a posteriori bound is the useful one in practice: **it estimates how far you are from arrival using only the size of your last step.** Worked at `delta = 0.12`, `Q_0 = 0`, drive `b = 0.3702` (the case of §44.8):

t	Q_t	true error	a priori bound	a posteriori bound
1	0.3702	2.7151	2.7151	2.7151
3	0.9828	2.1025	2.1026	2.1026
5	1.4571	1.6282	1.6282	1.6282
8	1.9757	1.1096	1.1096	1.1096

The bounds are attained with equality, to rounding. That is the affine case being exactly as strong as §44.0.3 said: distances shrink by exactly `1-delta`, so the inequality in Banach's bounds becomes an identity. In the nonlinear or multi-dimensional case the same bounds hold as inequalities and remain usable. **The practical instruction survives either way: your distance from arrival is about `(1-delta)/delta` times the size of your last step.** At `delta = 0.12` that multiplier is 7.33 — a step of 0.05 means you are roughly 0.37 from the fixed point.

44.0.7 The correspondence, discharged

Book One, Node 09 RECURSE	The term	Status
The spiral, not the circle	`Q_t = Q_inf + (1-delta)^t(Q_0 - Q_inf)`: the state changes every pass and never repeats exactly	**Proved** — §44.0.5
The winch post held fixed	`Q_inf` fixed across passes; only the transient moves	**Proved** — `Q_inf` has no `t` in it
The drift, checked weekly	`(1-delta)^t (Q_0 - Q_inf)`, the decaying inheritance	**Proved** — §44.0.5
The circle that does not climb	`eta = 0`: `b = -c*`, `Q_inf = -c*/delta < 0`	**Proved** — the circle is a regime whose fixed point is destruction

Book One, Node 10 FIXEDPOINT	The term	Status
The room exists	Banach existence, given `delta > 0` and completeness	**Proved** — §44.0.4
There is only one room	Banach uniqueness	**Proved** — §44.0.4
The room was always furnished	`Q_inf` independent of `Q_0`	**Proved** — §44.0.4
Arrival is recognition, not victory	geometric approach; most of the trajectory spent near `Q_inf`; the crossing is a JND event	**Derived** — §44.3, given a stated Weber fraction
Convergence is no certificate of worth	`eta F(u*) < c*` gives `Q_inf < 0` and the loop still converges	**Proved** — Theorem 3.5's own failure clause
The counterfeit arrival	a fixed point inflated by `x/delta` at another system's expense; sum invariant	**Derived** — §44.7

The two rows marked *derived* rather than *proved* are derived from the mathematics plus one badged postulate each, and both badges are shown where the derivation occurs. Nothing in this chapter reads a meaning into the theorem that the theorem does not carry.

44.0.8 Formula Validation Protocol — the error law

Per CANON_INTERNAL §6.

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`Q` in Q-units (the manual's unit of stored order; see Chapter 4 on conversion). `delta` dimensionless per period; `(1−delta)^t` dimensionless. Both sides of `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` are in Q-units. `t` in periods; the period must be declared.

Gate	Verdict	Content
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	The error law needs `Q` on an **interval** scale for differences and on a **ratio** scale for the percentage readings in §44.1. The manual's `Q` is postulated as a ratio-scale scalar with a true zero at "no stored order" (Appendix E's postulate set). **Badged: that `Q` is a scalar at all, rather than a vector of incommensurable stores, is a postulate of this manual and not a finding.** Where an operator's `Q` is measured by a Rasch-calibrated instrument (Chapter 54) the scale requirement is met; where it is measured by a self-rating out of 10, it is not, and only the ordering of successive gaps may be reported.
FVP-3 LIMITING CASES	PASS	`t -> inf`: gap `-> 0` for `delta in (0,1]`. `delta -> 0+`: gap `-> |Q_0 - Q_inf|`, no convergence; and `Q_inf -> inf` unless `b = 0`, which is the correct account of a system that never forgets and therefore never settles. `delta -> 1`: gap `-> 0` in one step; total forgetting means the state is entirely the current drive, `Q_t = b` for all `t >= 1`. `delta > 1`: excluded by the model's domain; `|1-delta| < 1` still holds up to `delta = 2`, giving oscillatory convergence, but `delta > 1` would mean losing more than is stored and the admissibility condition rules it out. `Q_0 = Q_inf`: gap is 0 for all `t` — you were already there.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* − c*)/delta`, unique by Banach. Meaning: **the state a life settles into under a given regime, computable in advance from four numbers, none of which is effort.** Second fixed point of the *error* dynamics: none — the error map `e -> (1−delta)e` has the unique fixed point `e = 0`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	From a single observed trajectory `{Q_t}` under constant drive, `delta` is identified from the decay of successive gaps, and the composite `b = eta M* − c*` is identified from `Q_inf * delta`. **`eta`, `M*` and `c*` are NOT separately identified from `Q` alone** — any triple with the same `eta M* − c*` produces an identical trajectory. Separating them requires independent measurement of at least two of the three: `M*` by magnitude estimation (Chapter 38), `c*` by a cost ledger (Chapter 5), `eta` then by residual. **State this whenever a trajectory is used to argue about efficiency rather than about outcome.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Paper E §5.3. Preregister the period, the `Q` instrument, and the window. Fit `delta` on a training block; **freeze**; forecast held-out `Q_t` from the frozen `delta` and the observed drive. Four baselines: persistence (`Q_{t+1} = Q_t`), direct-input (`Q_{t+1}` from `M_t` alone), matched AR(1), and a **shuffled-input control** in which the drive series is permuted. **The law dies if a frozen geometric-decay model does not beat the shuffled-input control on held-out data**, because that is the outcome in which the ordering of inputs — the entire content of "it accumulates" — carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `delta` recover across individuals, across investigators, across domains within a person? The manual does not know. §44.10 states the design that would answer it. Until it is run, a fitted `delta` is a within-person, within-domain descriptive parameter.

The Engineering

44.1 The convergence tables

Everything operational in this chapter comes off two constants. The **time constant**

```
tau = −1 / ln(1 − delta)        periods per factor of e
```

and the **iteration count to close a given fraction of the gap**

```
t(eps) = ln(eps) / ln(1 − delta)    periods to reduce the gap to a fraction eps
```

Both are exact; neither involves `Q_0`, `Q_inf`, `eta`, `M*`, or `C`. **How long arrival takes depends on the leak and on nothing else.**

delta	1−delta	tau	half-life	to 5%	to 1%	to 0.1%	to 0.01%
0.020	0.980	49.50	34.31	149	228	342	456
0.050	0.950	19.50	13.51	59	90	135	180
0.100	0.900	9.49	6.58	29	44	66	88
0.150	0.850	6.15	4.27	19	29	43	57
0.200	0.800	4.48	3.11	14	21	31	42
0.250	0.750	3.48	2.41	11	17	25	33
0.300	0.700	2.80	1.94	9	13	20	26
0.400	0.600	1.96	1.36	6	10	14	19
0.500	0.500	1.44	1.00	5	7	10	14
0.750	0.250	0.72	0.50	3	4	5	7

Read the table with the period attached, because the period is where the human meaning is.

A weekly practice with `delta = 0.05` per week is 59 weeks — fourteen months — from being within 5% of its destination. The same weekly leak restated per day is `delta = 0.0073`, needing 409 days: the same calendar. **A monthly cycle at `delta = 0.20` per month is fourteen months too.** The slightly deflating discovery in this table is that *changing cadence without changing the per-period leak changes nothing about the calendar*; cadence acts through `eta` and `C`, not through `tau`.

Two invariances that the table encodes and that are worth stating separately, because they are counter-intuitive and they are exact:

The cost of each halving is constant. At `delta = 0.10`, closing the gap from 100% to 50% takes 6.58 periods; from 50% to 25% takes 6.58 periods; from 1% to 0.5% takes 6.58 periods. The `ln(0.5)/ln(1−delta)` expression has no `t` in it.

The cost of each factor of ten is constant, and it is 21.85 periods at `delta = 0.10`.

interval	periods (at `delta = 0.10`)
close 0% -> 50% of the gap	7
close 50% -> 90%	15
close 90% -> 99%	22
close 99% -> 99.9%	22

The last ninety-nine hundredths of one percent costs exactly as much as the ninety to ninety-nine percent did, and more than the first half did. This is the arithmetic behind every experience of a practice that felt fast at the start and interminable at the end. The practice did not slow down. **The remaining distance is being measured in absolute terms by an instrument that works in ratios, and the ratios are being closed at a perfectly constant rate throughout.**

44.2 The shape of the approach

The fraction of the gap remaining, at `delta = 0.10`:

t	remaining	closed
0	1.00000	0.000%
1	0.90000	10.000%
2	0.81000	19.000%
5	0.59049	40.951%
10	0.34868	65.132%
20	0.12158	87.842%
29	0.04710	95.290%
44	0.00970	99.030%
66	0.00096	99.904%

A contraction spends the overwhelming majority of its life near its fixed point. At `delta = 0.10`, the system is within 5% of `Q_inf` from period 29 onward — which is to say that of the first 88 periods, 60 of them are spent inside a band of 5%. If the practice runs for two hundred periods, 171 of them are inside that band. **Whatever "being at the fixed point" feels like, it is what most of the trajectory feels like, and the transient is the unusual part.**

That fact is the entire explanation of §44.3.

44.3 Why arrival feels like recognition and not like victory

Four reasons, in increasing order of how much of the phenomenon they explain.

One: there is no arrival event. The gap is `(1−delta)^t` times its initial value, which is strictly positive for every finite `t`. **The system never reaches `Q_inf`. It approaches forever.** There is therefore no instant at which anything happens, and any felt instant of arrival is an event in the observer, not in the trajectory. Victory has a moment. This does not.

Two: the felt arrival is a threshold crossing in the observer's own instrument. If the operator can resolve a difference in `Q` of a fraction `c` — a Weber fraction, per Chapter 38 and §43.3 — then they notice arrival when the remaining gap falls below `c * Q_inf`, at

```
t_arrive = ln( c * Q_inf / |Q_0 − Q_inf| ) / ln(1 − delta)
```

For an operator starting from zero (`|Q_0 − Q_inf| = Q_inf`, so the ratio is just `c`):

delta c	0.02	0.05	0.10	0.20	0.30
0.05	77	59	45	32	24
0.10	38	29	22	16	12
0.20	18	14	11	8	6
0.30	11	9	7	5	4

Arrival is the period at which the remaining distance drops below your own just-noticeable difference, and it therefore happens at a different time for two people on identical trajectories. A coarse observer (`c = 0.30`) at `delta = 0.10` arrives at period 12; a fine one (`c =`

`0.02`) on the identical path arrives at period 38. Neither is wrong. They have different instruments. *(Badged: this uses Postulate R1's cousin — that Q is judged by the same log-register instrument as every other magnitude, hence a constant Weber fraction. That is an assumption of this manual, consistent with the psychophysics but not measured for Q.)*

Three: by the time you notice, you have been close for several time constants. From §44.2, an operator with `c = 0.05` and `delta = 0.10` "arrives" at period 29 — and has been within 10% since period 22, within 20% since period 16, within 35% since period 10. **The state that announces itself at period 29 has been substantially present since period 16.** The texture of recognition is the tail of an exponential, and the reason the destination seems familiar is that it has been the dominant feature of the last two-thirds of the journey.

Four — and this is the one that makes it recognition rather than mere familiarity: the destination never depended on you. `Q_inf = (eta M* − c*)/delta` contains no `Q_0`, no `t`, and no measure of effort. It is a property of the regime. From the first period, the trajectory was going exactly there. **You do not arrive somewhere new; you arrive where the parameters were always going to put you, and the recognition is accurate — the room really was furnished before you entered it, by the four numbers that describe how you live.**

That is the honest content of Book One's tenth node and it is not consoling. It says the same thing about a good regime and a bad one.

44.4 The invariant at arrival: headroom is pure geometry

A result that falls out of Theorem 3.5 and the headroom gauge together, and that this manual has not stated before.

Headroom (CANON §1.2, Paper E §6.4) is `h_t = eta ln N − delta Q_t − C_t`. Evaluate it at the fixed point. There, `delta Q_inf = eta M* − c*` by definition, so

```
h_inf = eta*lnN − delta*Q_inf − c*
      = eta*lnN − (eta*M* − c*) − c*

h_inf = eta * ( ln N − M* )                    THE ARRIVAL INVARIANT
```

Verified numerically across unrelated parameter sets:

eta	delta	N	u*	C	Q_inf	h computed from the gauge	eta(ln N − F)
0.35	0.12	9	0.55	0.22	3.0853	0.178789	0.178789
0.35	0.05	9	0.55	0.22	7.4048	0.178789	0.178789
0.20	0.30	6	0.30	0.10	0.2775	0.175094	0.175094
0.50	0.08	16	0.80	0.40	11.0309	0.103820	0.103820

At any fixed point, headroom equals the conversion efficiency times the surplus the geometry is failing to generate — and it does not depend on `delta` or on `C` at all. Rows one and two make the point: halving the leak more than doubles `Q_inf`, and leaves headroom at arrival unchanged to six decimal places.

The reading, and it is the most useful single sentence in this chapter for anyone deciding what to do next:

A converged system's remaining capacity for the new is exactly the surplus its geometry is not claiming. Not its slack, not its free time, not its energy. Its unclaimed u.

The table of that, at `eta = 0.35`, `N = 9`, `ln N = 2.1972`:

u*	N_eff	F(u*)	ln N − F	h_inf	% of ceiling claimed
0.00	1.000	0.0000	2.1972	0.7690	0.0%
0.25	3.000	1.0986	1.0986	0.3845	50.0%
0.40	4.200	1.4351	0.7621	0.2667	65.3%
0.55	5.400	1.6864	0.5108	0.1788	76.8%
0.70	6.600	1.8871	0.3102	0.1086	85.9%
0.85	7.800	2.0541	0.1431	0.0501	93.5%
1.00	9.000	2.1972	0.0000	0.0000	100.0%

A perfectly non-parallel life has zero headroom at arrival — it is already generating everything its category count permits, and the only remaining move is to raise N. That is Chapter 8's argument arriving from a completely different direction, and the convergence of the two is a point in the framework's favour, not a coincidence to be admired.

The dual reading is the harsh one. **An operator who feels enormous headroom at a settled state is not being told they have reserves. They are being told their framings are parallel.** h_inf is large precisely when F(u*) is small.

44.5 Where the leverage is: sensitivity of the fixed point

If arrival is determined by four numbers, the practical question is which of the four to move. Differentiate `Q_inf = (eta M* − c*)/delta`, at the base case `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55` (so `F = 1.6864`), `c* = 0.22`, giving `Q_inf = 3.0853`:

```
dQ_inf/d(eta)    =  F / delta                          =  14.0533
dQ_inf/d(C)      = −1 / delta                          =  −8.3333
dQ_inf/d(delta)  = −Q_inf / delta                      = −25.7111
dQ_inf/d(u)      =  eta(N−1) / ((1+(N−1)u) delta)      =   4.3210
```

Raw derivatives mislead because the levers are in different units. The elasticities — percent change in Q_inf per one percent change in the lever — are comparable:

Lever	Elasticity of Q_inf	Effect of a 10% improvement
eta (conversion)	**+1.594**	3.0853 -> 3.5772 (+15.94%)
delta (leak)	**-1.000**	3.0853 -> 3.4281 (+11.11%)
u (non-parallelism)	+0.770	3.0853 -> 3.3138 (+7.40%)
C (cost)	-0.594	3.0853 -> 3.2687 (+5.94%)

Three readings.

The elasticity with respect to `delta` is exactly -1, always. `Q_inf` is proportional to `1/delta`, so this is a structural fact and not a property of the base point: halve the leak, double the destination.

The elasticity with respect to `eta` is `eta*M*/(eta*M* - c*)`, which exceeds 1 whenever there is any cost at all, and diverges as the regime approaches unsustainability. At the base point it is 1.594. **A person operating near the sustainability boundary gets enormous returns from small improvements in how experience becomes character — and equally enormous losses from small degradations.** That is a real, derived, and slightly alarming property: fragility and leverage are the same number.

Cutting cost is the weakest of the four levers here, and it is usually the first one attempted. Elasticity -0.594 against `eta`'s +1.594: a 10% efficiency gain is worth 2.7 times a 10% cost cut at this base point. The ratio is `eta*M*/c*` and it favours efficiency whenever the regime is comfortably sustainable.

But there is a countervailing fact about `delta` that the elasticity hides, and §44.6 is it.

44.6 The cost of a long memory: noise amplification

Everything so far assumed constant drive. Real drive is not constant: `M_t` and `C_t` vary. The standard perturbation bound for a contraction: if the drive deviates from its mean by at most `epsilon` in each period, then in the long run

```
limsup_t | Q_t - Q_inf |  <=  epsilon / delta
```

The leak that determines how much you keep also determines how much of the input noise you keep. The amplification factor is `1/delta`, and it is the same `1/delta` that multiplies the signal.

`delta`	`1/delta`	wander at `eps=0.05`	at `eps=0.10`	at `eps=0.20`
0.020	50.00	2.500	5.000	10.000
0.050	20.00	1.000	2.000	4.000
0.100	10.00	0.500	1.000	2.000
0.150	6.67	0.333	0.667	1.333
0.200	5.00	0.250	0.500	1.000
0.300	3.33	0.167	0.333	0.667
0.500	2.00	0.100	0.200	0.400
0.750	1.33	0.067	0.133	0.267

The full trade, laid out on one line for three leak rates with the base case's drive:

`delta`	`Q_inf`	`tau`	periods to 5%	noise gain `1/delta`
0.06	6.1707	16.16	49	16.67
0.12	3.0853	7.82	24	8.33
0.24	1.5427	3.64	11	4.17

Halving the leak doubles the destination, doubles the time to reach it, and doubles the integrated noise. There is no free memory. A person with very low personal `delta` — who forgets nothing, whose skills do not atrophy, whose sense of self does not drift — accumulates a large `Q`, takes a very long time to change it, and carries every perturbation for twenty periods. That is not a defect of such a person; it is the same equation read on the other side.

The rule: **match `delta` to the volatility of the domain.** In a stable domain low `delta` is a gift; in a volatile one it is a liability, and deliberate shedding (Chapter 3; Book One's Node 11 SEAL) is the response.

44.7 Counterfeit arrival

Book One's tenth node warned that a well-converged state can still be the wrong one, since convergence is a property of loops and not a certificate of worth. Theorem 3.5 already supplies the first version of that: **if `eta F(u*) < c*` the fixed point is negative, and the loop converges anyway — toward the destruction of the thing.** Paper E does not hide this and neither does this manual.

This section supplies the second version, which is the one the manifest names: **a fixed point reached at another's expense — the room furnished with someone else's goods.**

44.7.1 The transfer model Two systems, `A` and `B`, each running the storage law, with a transfer `x >= 0` per period from `B` to `A`. `A` receives `x` as a free addition to its drive; `B` pays it as an addition to its cost.

```
Q_A(t+1) = (1−delta) Q_A(t) + eta∗F_A − C_A + x
Q_B(t+1) = (1−delta) Q_B(t) + eta∗F_B − C_B − x
```

Both remain affine contractions, so both still converge, and the fixed points are

```
Q_A,inf(x) = ( eta∗F_A − C_A + x ) / delta  =  Q_A,inf(0) + x/delta
Q_B,inf(x) = ( eta∗F_B − C_B − x ) / delta  =  Q_B,inf(0) − x/delta
```

The transfer moves each fixed point by exactly `x/delta`, in opposite directions, and the sum is unchanged. Numerically, with `eta = 0.35`, `delta = 0.12`, `C = 0.22`, `N = 9`, `u_A = 0.55` (`F_A = 1.6864`), `u_B = 0.45` (`F_B = 1.5261`):

x	Q_A,inf	gain to A	Q_B,inf	sum
0.0000	3.0853	—	2.6177	5.7030
0.1000	3.9187	+0.8333	1.7843	5.7030
0.2000	4.7520	+1.6667	0.9510	5.7030
0.3141	5.7028	+2.6175	0.0002	5.7030
0.4000	6.4187	+3.3333	-0.7157	5.7030

The sum column is the finding. Transfer relocates a fixed point; it creates none. No amount of extraction increases the total stored order of the pair by one unit. This is not an ethical claim dressed as arithmetic — it is arithmetic, and its ethical reading is left to the reader, who already has CANON §5.4 in hand.

44.7.2 The exhaustion horizon There is a critical transfer rate at which B's fixed point hits zero:

```
x* = eta * F_B − C_B
```

which for the case above is `x* = 0.35 * 1.5261 − 0.22 = 0.3141`. This is simply Corollary 3.4 — the sustainability ceiling — read as a constraint on how much can be taken. **Below x*, B is diminished but survives at a lower fixed point. Above x*, B's fixed point is negative and B runs out.**

Take `x = 0.40 > x*`. B's fixed point is `−0.7157`. Starting from its undisturbed fixed point of 2.6177, B's trajectory is:

t	Q_B
0	2.6177
1	2.2177
3	1.5559
5	1.0434
7	0.6466
9	0.3393
11	0.1013
12	0.0032
13	**-0.0830**

The exhaustion horizon in closed form, from the error law:

```
t_exhaust = ln( −Q_B,inf / (Q_B,0 − Q_B,inf) ) / ln(1 − delta)  = 12.04 periods
```

Twelve periods. During all twelve, A is converging comfortably toward 6.4187 — a fixed point 108% higher than its own regime could produce — and every measurement A can take of itself says the practice is working.

44.7.3 What happens at the end When B fails, the transfer stops. A's drive reverts and its fixed point drops back to 3.0853. A now has a gap of 3.3333 Q-units to close, downward, and it closes it geometrically at the same `delta`: **24 periods to come within 5% of the new, lower destination.**

```
A's peak fixed point during extraction : 6.4187
A's fixed point after B fails          : 3.0853
loss on collapse                    : 3.3333 Q-units  (51.9% of the peak)
twelve periods of gain, twenty-four of unwinding
```

The counterfeit arrival is not merely unjust; it is arithmetically temporary, and its unwind is twice as long as its build. That asymmetry is not moralising — it is the ratio of the exhaustion horizon to the reconvergence time, and both were computed from the same `delta`.

44.7.4 The test that distinguishes genuine from counterfeit arrival Here is the operational point of the whole section, and it is one line.

An arrival is genuine if every party to it has `h_t >= 0` under its own regime, counting the transfer as a cost. It is counterfeit if any party has `h_t < 0`.

That is Corollary 3.4 used as an audit rather than as a forecast. The check is cheap. For each party, compute `h = eta ln N − delta Q − C_total` where `C_total` includes every transfer that party is paying. If any h is negative, then by Corollary 3.4 **no choice of interaction geometry prevents that party's order from decreasing** — and the arrangement's apparent stability is being paid for out of a store that is running down.

Two honest complications.

The audit requires the other party's ledger. `A` cannot run it from inside `A`; it needs `B`'s `Q`, `C`, and `delta`. **This is precisely why the check is skipped, and precisely why it is the check** — at personal scope, it is the arrangement whose sustainability you can only assess by asking a question you would rather not ask.

A negative h in a party who has not been asked is not established. CANON §5.4 governs. The audit is run by request, with a decline path, and a party who declines to supply their ledger is recorded as unknown, not as consenting. **An audit that conscripts the other person's data to reassure the person doing the extracting has reproduced the problem it was meant to detect.**

44.8 Arrival, stall, and the limit of what arithmetic can tell you

A flat `Q` has exactly one arithmetic cause: `eta M_t = delta Q_t + C_t`, the one-step growth condition of Corollary 3.2 holding with equality. That single equation covers both of the following:

- **Arrival.** The practice has converged to a fixed point that its operator chose, by choosing the regime that generates it.
- **Stall.** The practice has converged to a fixed point nobody chose, inherited from parameters nobody set, and the operator has been calling it stability.

The arithmetic cannot distinguish these. They are the same equation. This manual states that plainly rather than manufacturing a discriminator, because the honest position is that the difference is not in the trajectory — it is in the record of how the parameters came to have their values.

What *can* be checked, and what the specification below requires:

1. **Was `Q_inf` computed before it was reached?** A chosen fixed point is one that was predicted. Chapter 24's survey and the four parameters give a forecast; a forecast recorded in advance and then met is evidence of a chosen regime.
2. **Is `h_inf = eta(ln N − F(u*))` known and accepted?** An operator at a chosen fixed point can state how much of the ceiling their geometry is claiming, and has a reason for the remainder.
3. **Is there a sunset?** Book One's Node 11 SEAL and Chapter 23's shedding: a chosen regime has a review date. A stall has never had one.

Arrival and stall are distinguished by documentation, not by measurement, and that is not a weakness of the model — it is the model correctly declining to certify an intention it cannot observe.

44.9 Changing where you arrive

The consolidated instruction, from §44.4, §44.5, and §44.6 together.

```
 to move the destination      change eta, u (through F), C, or delta
to move faster to it      change delta -- and accept that this moves the destination too
to trade destination      raise eta or u; these move Q_inf without moving tau
   without changing speed
to raise the ceiling      raise N (Chapter 8); this is the only move that raises
                               the maximum any geometry could reach
```

And the negative result: **trying harder inside a fixed regime does not change where you arrive.** Effort that moves none of the four parameters moves only the transient — the same path, the same room, the same rate, since the rate is `delta` and effort is not `delta`. The sensation of pushing is real and it is spent on the term that was going to decay anyway.

44.10 What would need measuring

Per CANON §2, the agenda rather than the results:

1. **A `Q` instrument with ratio-scale properties.** Rasch-calibrated where possible (Chapter 54), because specific objectivity is exactly the property that lets a `Q` measured with one item set be compared with a `Q` measured with another. Without it, FVP-2 stays conditional.
2. **`delta` by domain and by person.** Fit from the decay of successive gaps under a constant regime, per FVP-6, with the freeze step enforced. Report distributions, not point estimates.
3. **The Weber fraction for `Q`.** This determines `t_arrive` in §44.3 and it is currently a parameter of a table rather than a measured constant.
4. **The joint identifiability study.** Independent measurement of `M*` (magnitude estimation) and `c*` (cost ledger) in the same subjects, to test whether `eta` recovered by residual is stable — which is FVP-5 and FVP-7 attacked together.

The Specification

CS-1 — Declare the period. Every `delta`, `tau`, and iteration count is meaningless without the period attached. The period is fixed before data collection and is not changed mid-block.

CS-2 — Compute `Q_inf` before pursuing it. `Q_inf = (eta M* − c*)/delta` from the four measured parameters, recorded in writing, before the practice is undertaken. *(This is the requirement that makes §44.8's first check possible at all.)*

CS-3 — Verify `delta > 0`. If the practice has no loss term — nothing is forgotten, nothing atrophies, no maintenance is required — there is no fixed point and the convergence tables do not apply. State the loss mechanism explicitly.

CS-4 — Check the sign of the fixed point. If `eta M* < c*` then `Q_inf < 0` and the regime is unsustainable. **Report this before proceeding; do not run the practice and discover it.** *(Theorem 3.5's own failure clause.)*

CS-5 — State the arrival criterion in advance. Declare `c`, the fraction of `Q_inf` at which the practice will be called arrived, and compute `t_arrive` from §44.3. An undeclared arrival criterion converts an asymptote into an argument.

CS-6 — Use the a posteriori bound to track progress. Distance from arrival is approximately

`((1−delta)/delta)` times the last observed step. Report it each period; it costs one multiplication.

CS-7 — Compute `h_inf = eta(ln N − F(u*))` at the fixed point and account for it. Large headroom at arrival is a statement about parallel framings, not about reserves. *(§44.4.)*

CS-8 — Match `delta` to domain volatility. Where the domain changes faster than `1/delta` periods, deliberate shedding is specified, not optional. *(§44.6.)*

CS-9 — Run the transfer audit. For any arrangement in which one party's drive includes a transfer from another, compute h for **every** party with the transfer counted as that party's cost. Record any negative h. *(§44.7.4.)*

CS-10 — The audit has a decline path. No party is required to disclose their ledger. A declined audit is recorded as *unknown*, never as *satisfied*. *(CANON §5.4, binding.)*

CS-11 — Do not compare `delta` across people without recalibration. FVP-7 is open. A fitted `delta` is within-person and within-domain until the reproducibility study exists.

CS-12 — Freeze before claiming. FVP-6 in full, including the shuffled-input control.

The Failure Modes

Fault	Diagnostic signature	Fix
The unarrived asymptote	Operator reports "almost there" for many periods and treats it as failure	There is no arrival event (§44.3). Declare `c` in advance per CS-5 and use it. The last 0.9% costs what the previous 9% did
Effort substitution	Increasing effort with no change in `eta`, `u`, `C`, or `delta`; `Q` follows the same path	Effort acts on the transient, not the fixed point (§44.9). Identify which of the four parameters the effort is supposed to move; if none, stop
The negative fixed point run anyway	`eta M* < c*`; `Q` declining steadily; practice continued because "consistency matters"	CS-4. The loop converges toward destruction. Cut `C`, raise `eta` or `u`, or stop
Zero-delta fantasy	A plan with no loss term: nothing decays, nothing needs maintenance	No contraction, no fixed point, unbounded linear growth in the model and rapid divergence from reality. Name the loss mechanism (CS-3)
Low-delta noise integration	Very stable person or practice; large swings in `Q` traced to small perturbations months earlier	`1/delta` amplification (§44.6). Either accept it as the price of memory or raise `delta` deliberately by shedding

Fault	Diagnostic signature	Fix
Headroom misread as reserve	Large h at a settled state, read as "capacity available"	`h_inf = eta(ln N − F(u∗))`. Large h means parallel framings (§44.4). Raise u or N, do not "use the capacity"
Counterfeit arrival	A party's Q rising steadily with no change in their own `eta`, u, `C`, or `delta`	The extra drive is a transfer. Run CS-9 on every party. Sum is invariant (§44.7.1)
Extraction above x∗	Counterparty's Q declining linearly toward zero on a finite horizon	Compute `x∗ = eta F_B − C_B` and `t_exhaust`. Above x∗ the arrangement has a countdown, and the unwind is longer than the build
Stall named arrival	Flat Q, no predicted `Q_inf` on record, no review date	§44.8. The arithmetic cannot adjudicate; the documentation can. If there is no forecast and no sunset, treat as stall
Cadence confusion	Cadence increased, calendar unchanged, operator frustrated	`tau` depends on the per-period `delta`. Changing cadence without changing per-period leak changes nothing (§44.1). Cadence acts through `eta` and `C`
Cross-person `delta` comparison	"Her `delta` is lower than mine"	FVP-7 is untested. CS-11
Ordinal Q	Convergence rates fitted to a 1–10 self-rating	FVP-2. Differences on an ordinal scale are not differences. Report the ordering of successive gaps only, or build a calibrated instrument

---

Worked Example

A practice rebuilt after a move. As in Chapter 43, the data are **illustrative** — a constructed case used to demonstrate the computation. No claim is made that these parameters describe any real person.

Setting. An operator restarting a daily practice from zero after a relocation that reset their stored order. Period = one week. They measure, per the protocols of Chapters 38, 39, 40 and 5:

```
N     = 9       distinguishable framings available in the domain
u∗    = 0.55    measured non-parallelism among their own framings
eta   = 0.35    conversion of moment into durable structure
delta = 0.12    weekly loss without maintenance
```

```
c*    = 0.22     weekly maintenance cost, from the cost ledger
Q_0   = 0        the move reset the store
```

Step 1 — the ceiling and the drive

```
ln N     = ln 9          = 2.1972   the hard ceiling on surplus per interaction
N_eff     = 1 + (9-1)(0.55)   = 5.400     an audited headcount of framings
F(u*)     = ln(5.400)        = 1.6864   the surplus this geometry generates
claimed    = 1.6864 / 2.1972   = 76.8%     of the ceiling
drive b    = 0.35*1.6864 - 0.22 = 0.3702   Q-units per week
```

b > 0, so CS-4 passes: the regime is sustainable and the fixed point is positive.

Step 2 — the destination, computed before departure (CS-2)

```
Q_inf = ( 0.35 * 1.6864 - 0.22 ) / 0.12 = 3.0853 Q-units

upper bound check (Thm 3.3, Q_0 = 0):  eta*lnN/delta = 6.4086
3.0853 <= 6.4086.  Consistent.
```

The room is furnished at 3.0853 Q-units, and this was known in week zero.

Step 3 — the schedule

```
tau        = -1/ln(0.88)          =  7.82 weeks
half-life  = ln(0.5)/ln(0.88)     =  5.42 weeks
to 5%      = ln(0.05)/ln(0.88)    = 24 weeks
to 2%                             = 31 weeks
to 1%                             = 37 weeks
```

Declaring c = 0.05 per CS-5: **arrival is week 24, and this is written down in week zero.**

Step 4 — the trajectory

t (wk)	Q_t	gap	gap / Q_inf	h_t	step dQ
0	0.0000	3.0853	1.0000	0.5490	0.3702
1	0.3702	2.7151	0.8800	0.5046	0.3258
2	0.6961	2.3893	0.7744	0.4655	0.2867
4	1.2351	1.8503	0.5997	0.4008	0.2220
6	1.6525	1.4328	0.4644	0.3507	0.1719
8	1.9757	1.1096	0.3596	0.3119	0.1332
10	2.2261	0.8593	0.2785	0.2819	0.1031
12	2.4199	0.6654	0.2157	0.2586	0.0799
14	2.5700	0.5153	0.1670	0.2406	0.0618
16	2.6863	0.3990	0.1293	0.2267	0.0479
20	2.8460	0.2393	0.0776	0.2075	0.0287
25	2.9590	0.1263	0.0409	0.1939	0.0152
30	3.0187	0.0666	0.0216	0.1868	0.0080

Closed-form check at `t = 30`: `Q_inf + (1−delta)^30 (0 − Q_inf) = 3.018683`. The iteration and the error law agree to six figures, as §44.0.5 requires.

Step 5 — what the operator experiences, week by week

Weeks 1–5. Steps of 0.37, 0.33, 0.29, 0.25, 0.22, closing 47% of the gap. This is the part everyone calls "fast at the beginning" — it is not fast; it is the same 12% per week as every other week, applied to a large number.

Weeks 6–16. Steps from 0.17 down to 0.048. The gap goes from 46% to 13%. This is the long middle, and per §44.1 it costs what it costs: closing the gap from 50% to 10% takes 12.59 weeks here, against the 5.42 weeks the first halving took. **Nothing has gone wrong. The rate has not changed and cannot change.**

Weeks 17–24. Steps below 0.045. The a posteriori bound (CS-6) at week 20: the step just completed, week 19 to week 20, was 0.0326, giving distance from arrival `<= (0.88/0.12) * 0.0326 = 0.2393` — against a true gap of 0.2393. The bound is attained exactly, for the reason §44.0.6 gives. **The operator can compute their remaining distance from the size of their last week alone, without knowing `Q_inf` at all.**

Week 24. Declared arrival. Gap is under 5% of `Q_inf`. It does not feel like an event, because it is not one — the state at week 24 differs from the state at week 20 by 0.0958 Q-units, about 3.1% of the destination, spread over four weeks, which is at or below most plausible resolutions. **The operator recognises the state because they have been substantially in it since week 16.**

Step 6 — the headroom reading at arrival

```
h_inf = eta ( ln N − F(u*) ) = 0.35 * ( 2.1972 − 1.6864 ) = 0.1788
```

Cross-check from the gauge at `Q_inf`: `0.35*2.1972 − 0.12*3.0853 − 0.22 = 0.178789`. Identical, as §44.4 requires.

The operator arrives with 0.1788 Q-units per week of unclaimed capacity, and §44.4 says exactly what that number is: it is the 23.2% of the ceiling their framings are not generating. It is not free time and it is not energy. **The only way to convert it is to raise `u` toward 1 or to raise `N` above 9,** and Chapters 8 and 39 are where those moves live.

Step 7 — the counterfactual, and the levers

Suppose at week 24 the operator wants a higher destination. From §44.5:

Change	New `Q_inf`	Change
baseline	3.0853	—
`eta` 0.35 -> 0.385 (+10%)	3.5772	+15.94%
`delta` 0.12 -> 0.108 (-10%)	3.4281	+11.11%
`u` 0.55 -> 0.605 (+10%)	3.3138	+7.40%
`C` 0.22 -> 0.198 (-10%)	3.2687	+5.94%

The 10% efficiency improvement is worth 2.7 times the 10% cost cut. And the `delta` move carries the §44.6 penalty: at `delta = 0.108` the time constant rises from 7.82 to 8.75 weeks, the time to 5% from 24 to 27 weeks, and the noise gain from 8.33 to 9.26. **Reducing the leak buys a better room, more slowly, with more of every bad week retained.**

Step 8 — the audit that must not be skipped

CS-9. The operator's drive is `eta F(u*) = 0.5902` per week and every component of it is their own: their framings, their conversion, their cost. There is no `x`. The audit passes trivially.

Had there been a transfer — a household absorbing the practice's cost, a colleague covering the work it displaces — the audit would require **that party's** `h`, computed with the transfer as their cost, obtained by asking, with a decline path (CS-10). Per §44.7, an extraction of `x = 0.40` per week against a counterparty whose sustainable ceiling is `x* = 0.3141` would give that party a fixed point of -0.7157 and an exhaustion horizon of **12.04 weeks** — during which the operator's own numbers would look excellent and their fixed point would read 6.4187 instead of 3.0853.

Every measurement the operator could take of themselves would say the practice was working twice as well as it does. That is the whole danger of counterfeit arrival: it is not detectable from inside the system that benefits from it. Only the other ledger shows it, and only if it is asked for.

The one line to carry out of the chapter

```
Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )
```

Where you are going does not contain where you started. Where you are is where you are going, plus what is left of where you started, and what is left of where you started is halving every 5.42 weeks. Arrival is not an achievement and not a destination reached by force. It is the moment the second term drops below the resolution of the instrument reading it — and the room, by then, has been furnished for a long time.

Chapter 45 . The Personal Formulary

The Law

This chapter is a lookup table. It is the payoff of Part VIII: every formula that applies at the **internal scope** — psychology, personal, one mind — collected in one place, in one format, with one unit system, each carrying its full Formula Validation Protocol so that a reader can tell in ten seconds which gates it passes and which it does not.

A formulary is not a theory. It is a shelf. Its virtue is that the entries are in the same units, formatted identically, and honest about their limits, so that a person holding a real problem can find the right instrument without reading a book. The chapters before this one made the arguments; this one exists so the arguments do not have to be re-read.

45.0.1 How to use it

Three access paths. **By quantity:** the index in §45.0.2 maps each parameter (u, eta, h, r) to its entry. **By fault:** every entry ends with a fault table keyed to the felt complaint, and *The Failure Modes* below covers the faults that come from using the collection rather than any single entry. **By sequence:** *The Specification* below gives the dependency order for a full review. **Running the entries out of order produces numbers that look fine and mean nothing, because six of the nine depend on quantities the others produce.**

45.0.2 The nine entries

#	Name	Formula	Primary parameter	Source
F-01	The Moment	`M = (D/P) . E`	`M`	Curzi; Ch 34
F-02	The Calibration Ratio	`r = D/P, Delta = |ln r|`	`r`	CANON_INTERNAL §4; Ch 35
F-03	The Internal Surplus	`F_int = ln(1 + (N−1)u)`	`u, N_eff`	Paper A Thm 2.1; Ch 39
F-04	The Self-Storage Law	`Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`	`Q, eta, delta`	Paper E §3; Ch 40
F-05	Personal Headroom	`h_t = eta ln N − delta Q_t − C_t`	`h`	Paper E §6.4; Ch 6, Ch 40
F-06	The Launch Budget	`L = max(0, M − C)`	`L`	CANON_INTERNAL §8; Ch 41
F-07	The Gradient	`g = −grad C`	`g`	Book One Node 02; Ch 42
F-08	The Return Difference	`Delta_R = |ln(M'/M)|`	`rho`	CANON_INTERNAL §8; Ch 43

#	Name	Formula	Primary parameter	Source
F-09	The Contraction	`|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`	`Q_inf`, `tau`	Paper E Thm 3.5; Ch 44

Four of the nine are **theorems** of the corpus (F-03, F-04, F-05, F-09) and appear here in their internal-scope reading. Two are **the author's formula** and its immediate consequence (F-01, F-02). Three are **derived formulas of this manual** (F-06, F-07, F-08), derived in the sense CANON_INTERNAL §8 requires and badged where their derivations pass through a postulate.

Every entry states which of the three it is in its first line, because the difference decides how much weight it can carry.

45.0.3 The entry format

Every entry has the same eight parts, in the same order, so the chapter can be read as a table:

```
  STATEMENT          the formula, in a fence, with every symbol defined
  READS              one sentence: what the formula says in words
  STATUS             theorem / author's formula / derived-with-badge
  DERIVATION         where it comes from in the corpus, and what it needs
FVP           all seven gates, verdict plus content        (CANON_INTERNAL §6)
  PROTOCOL           numbered measurement procedure
  WORKED NUMBER      one computed instance
  FAULTS             named failure modes with signature and fix
```

45.0.4 The unit system

Everything internal is measured in **nats**, or in **Q-units** (the manual's unit of stored order, defined in Chapter 4), or in **Q-units per period**. Because the four registers converge (CANON_INTERNAL §1), any quantity in nats is immediately readable in three other registers by exact constants:

```
  1 nat    = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
  1 bit    = 0.693147 nats = 2 octaves EXACTLY = 2400 cents exactly
  1 octave = 0.346574 nats = 0.5 bits          = 1200 cents

surplus -> frequency:  f = k * e^(2F),  k = 0.483425732988414 Hz  (CANON_SONIC)
```

The ceilings, which bound five of the nine entries:

N	ln N (nats)	bits	octaves	f = k e^(2 ln N) (Hz)
2	0.6931	1.0000	2.000	1.9337
4	1.3863	2.0000	4.000	7.7348
6	1.7918	2.5850	5.170	17.4033
7	1.9459	2.8074	5.615	23.6879

N	ln N (nats)	bits	octaves	f = k e^(2 ln N) (Hz)
9	2.1972	3.1699	6.340	39.1575
10	2.3026	3.3219	6.644	48.3426
16	2.7726	4.0000	8.000	123.7570

`ln N` is the hard ceiling on surplus per interaction (Paper D, Corollary 2.2), and by the forced consequence of CANON_INTERNAL §1 it is also the ceiling on the moment: the richest available moment is bounded by the log of how many kinds of thing you can distinguish. That is Ashby's requisite variety appearing in first-person experience. It was not chosen; it follows.

45.0.5 The circuit — how the nine connect

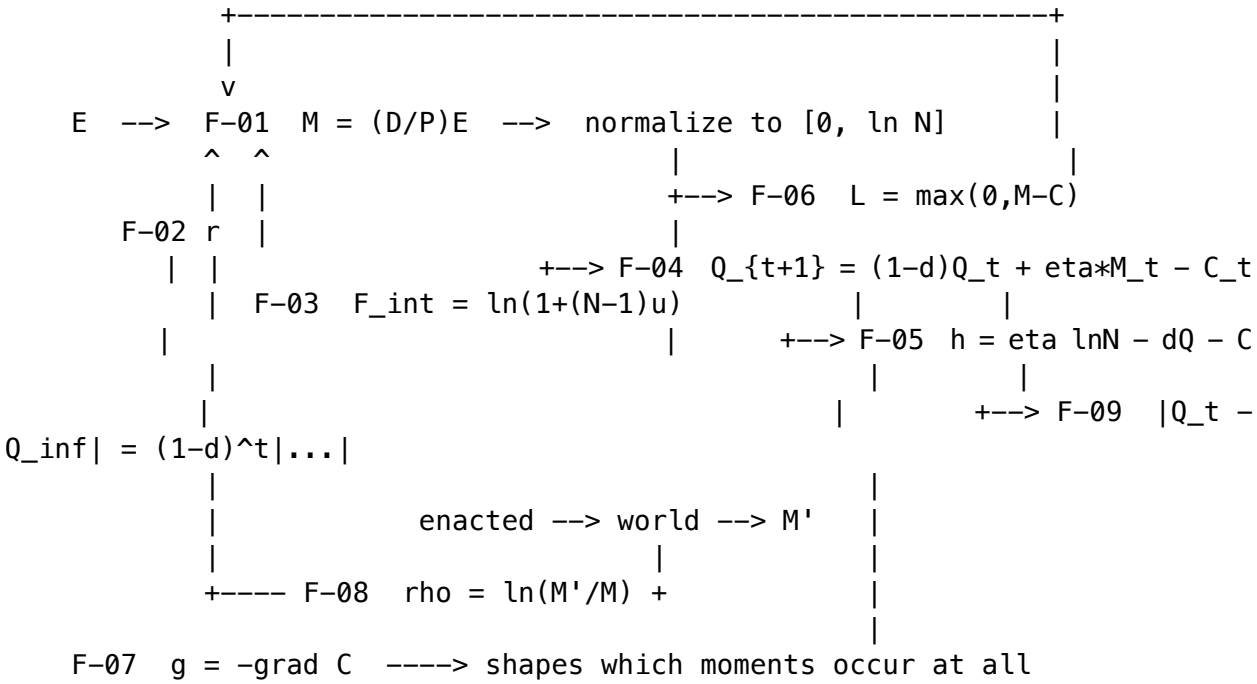
E --> F-01 M = (D/P)E --> normalize to [0, ln N]
F-02 r
F-03 F_int = ln(1+(N-1)u)
F-04 Q_{t+1} = (1-d)Q_t + eta*M_t - C_t
F-05 h = eta lnN - dQ - C
F-06 L = max(0,M-C)
F-08 rho = ln(M'/M) +
F-09 |Q_t - Q_inf| = (1-d)^t|...|
enacted --> world --> M'
F-07 g = -grad C ----> shapes which moments occur at all

Three structural facts follow, and they govern the order of operations in *The Specification*.

F-03 sets the ceiling F-01 lives under — no moment is richer than `ln N`, and `N_eff` says how much of `N` is real. **F-01 is the only producer:** storage, headroom, launch and arrival are all fed by `M`, so **if the Moment is mismeasured, all six downstream numbers are wrong in the same direction and will look consistent with one another, which is the most dangerous kind of wrong. F-08 is the only feedback path, and it acts on `D`, not on `Q`** — the echo corrects the frame, which changes the integrator's input, not its state, which is why feedback improves character slowly (Chapter 43, §43.10).

45.0.6 The FVP verdict legend

CANON_INTERNAL §6 defines the seven gates. This chapter uses six verdict words and they are not interchangeable:

Verdict	Meaning
PASS	The gate is satisfied as the formula is stated here
CONDITIONAL PASS	Satisfied only under a stated condition; if the condition fails, the formula is **demoted in print** to a heuristic ordering
DEMOTE	The gate is **not** satisfied as things stand. The formula's structure — limits, fixed points, orderings — remains in force; its magnitudes are withdrawn as measured quantities until the named instrument exists
PARTIAL	Satisfied for some parameters and not others; the unrecoverable ones are named
SPECIFIED	Not yet run, but the exact procedure that would run it is given here
UNTESTED — declared	Not run, no data, and the manual says so rather than implying otherwise

Nine entries, sixty-three gates, and not one is marked PASS on the reproducibility gate, because this manual has run no cross-individual study of any of these quantities. CANON_INTERNAL §6 names FVP-7 as the gate that decides the hard-science question and the one most often skipped; a formulary's honest option is to leave it visibly open on every page rather than quietly on none.

45.0.6.1 Two standing verdicts that govern more than one entry Two findings arrived after the entries below were drafted, and both cut across several of them. They are stated once here, in force everywhere, rather than repeated nine times.

STANDING VERDICT 1 — every `M`-derived entry is DEMOTED on gate 2. Chapters 52 and 53 run the Formula Validation Protocol on the Moment Equation and return **DEMOTE on FVP-2**: `D` and `P` are not instrumented on a common ratio scale with a shared true zero, and division without a true zero yields a number with no referent. This governs **F-01, F-02, F-04, F-06 and F-08** — every entry whose value passes through `M` or through `r = D/P`.

> **NOT REPORTABLE AS MEASURED QUANTITIES:** the magnitude of `D/P`; the distortion `Delta = |ln r|` as a number; the quantitative form of the `M –> Q` join (F-04). Every such figure in the worked numbers below is an **illustrative construction demonstrating the arithmetic**, and is labelled as such at its use.
>
> **STILL IN FORCE:** the limiting cases; the `D = P` fixed point; and the ordering *over-framed* / *calibrated* / *under-framed*. **These require only ordinal information and the demotion does not touch them.** What is demoted is magnitude, not structure.

The route out is constructive, specific and priced. Chapter 54, **Theorem 54.3**, given badged Postulate SO1: measure `D` and `P` as Rasch person-measures on one common linked scale, and `Delta = |ln(D/P)| = |theta_D − theta_P|` becomes genuinely ratio-scaled — Rasch separation supplies the true zero that division requires. The cost, from §54.12: **approximately 123 items per term to establish a half-logit distortion at 95 per cent**, and about 31 items for a

full logit. A twelve-item pair of scales resolves nothing finer than roughly 1.6 logits. **A formulary that demotes five of its own nine entries, names exactly what survives, and prices the study that would restore the rest is the protocol working. It is not a caveat and it is not an apology.**

STANDING VERDICT 2 — the shuffled-input control has three measured failure modes, and every FVP-6 row below inherits all three. The four-baseline suite of Paper E §5.3 is cited in every falsification row in this chapter. Chapters 55 and 81 measured what it actually catches, by simulation rather than assertion:

#	Failure mode	Measured
1	**Anti-conservative as ordinarily run.** Against a nominal `alpha = 0.05` across 500 null trials, iid permutation false-rejected at **0.108 — 2.2x nominal**; circular shift at 0.120; block permutation at 0.060, near nominal	Ch 55
2	**Blind to repackaging.** A filtered repackaging of the input passed all four baselines including the shuffle at `p = 0.0005`, and was caught **only by the structural ceiling test** (Paper E Cor 4.2) — its single-step jump was **38.5x** the theoretical maximum	Ch 25, Ch 55
3	**Blind to shared slow drivers.** In 4,000 trials the shuffle reported honestly in the clean case (48.17% detection) and **passed 98.35% of the time against a confound driving both series**	Ch 79–81

> **THE RULE, AND IT SUPERSEDES ANY EARLIER FORM: the shuffle is a weak, mis-calibrated empirical audit. Repackaging is caught only by the structural ceiling test; confounding is caught only by design controls — blinding, channel separation, prospective sampling, a yoked control. No surrogate catches either.** Calibrate the surrogate by simulation on your own design before you use it, and **design controls, not more surrogates.** Chapter 55 specifies both.

45.0.7 The refusals, as they apply to a formulary

CANON §5 governs every entry. Three of the four refusals have specific formulary consequences:

No unbadged postulates — every entry whose derivation passes through an assumption states

it in the entry. **No uncomputed numbers** — every worked number here was computed, from data that are **illustrative constructions**, stated as such at each use. **No intervention that removes the ability to decline** — at internal scope other people appear in F-08's echo sources and in F-05's cost ledger, and both are governed: **an echo source may decline and is recorded as missing data, never as a null return; and a cost carried by someone else is counted in their ledger, by asking, with a decline path.**

The Engineering

F-01 . THE MOMENT

STATEMENT

```
M = ( D / P ) . E

D  perspective   the frame brought to the encounter
P  perception    the intake taken from the encounter
E  environment   what is there
M  the moment    experience as constructed

r = D/P must be dimensionless, so D and P must share units.
M then inherits E's units.
Normalized form for use downstream:  M_norm in [0, ln N]
```

READS. *The moment is the environment, scaled by how far your frame outruns or falls short of your intake.*

STATUS. The author's formula (Curzi), analysed in full in Chapter 34. Not a theorem of Papers A–E; it is the object Volume Two is built to make measurable.

DERIVATION. `M` is not derived from the surplus law — it is the corpus's second formal object, joined to the first through the storage law (F-04), where `M` supplies the input `S_t`: *perspective and perception produce the moment; the moment becomes character* (CANON_INTERNAL §8). The ceiling `M <= ln N` is not assumed but forced, once `M` is used as `S_t`, by `0 <= S_t <= ln N` (Paper D, Cor 2.2).

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`D` and `P` in shared units (whatever the estimation procedure supplies); `r` dimensionless; `M` in `E`'s units. For downstream use, `M` is normalized into nats on `[0, ln N]`, and the normalization must be stated.

Gate	Verdict	Content
FVP-2 SCALE TYPE	**DEMOTE — the crux** (Standing Verdict 1)	Division requires ratio scale with a true zero for `D` and `P`, on a **common** scale. No such instrument yet exists, so **the magnitude of `D/P` and every quantity derived from it are withdrawn as measured quantities (Ch 52, Ch 53).** `M = (D/P)E` reads as the ordering *over-framed* / *calibrated* / *under-framed*, plus its limits and fixed point — all of which need only ordinal information and none of which is affected. **Route out: Ch 54 Thm 54.3**, Rasch logit difference given badged Postulate SO1, at ~123 items per term for half a logit.
FVP-3 LIMITING CASES	PASS	`P -> 0`: `M -> inf` — frame with no intake: hallucination, dogma, apophenia. `D -> 0`: `M -> 0` — intake with no frame; **sensation is not experience**. `E -> 0`: `M -> 0` — no environment, no moment. Each divergence names a real phenomenon, which is what the gate demands.
FVP-4 FIXED POINTS	PASS	`D = P` gives `M = E`: **undistorted contact.** The formula's most important content — when frame matches intake, the moment equals the environment.

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL**	`M` and `r` are recoverable. `D` and `P` **individually are not identified from `M` alone** — any pair with the same ratio and the same product-with-`E` predicts identically. Recovering them separately requires manipulating one while holding the other: change the frame with the environment fixed, or change the environment with the frame fixed.
FVP-6 FALSIFICATION	SPECIFIED	Paper E §5.3. Preregister the estimation procedure for `D`, `P`, `E`. Fit on a training block, freeze, predict held-out `M`. Beat persistence, direct-input (`E` alone), matched AR(1), and the shuffled-input control. **Killed if `E` alone predicts `M` as well as `(D/P)E` does** — that outcome says the frame ratio carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Cross-individual recovery of `D`, `P`, `E` under a common operationalization has not been run. Chapter 54's Rasch approach — separating person parameters from item parameters — is the route, and until it exists `M` is a within-person quantity.

PROTOCOL

1. Fix a **modulus**: assign a magnitude (commonly 50) to one reference encounter and hold it for the whole block.
2. **Before** the encounter, estimate `D`: how much structure you are bringing — expectation, model, prior commitment — against the modulus.
3. **During or immediately after**, estimate `P`: how much you actually took in.
4. Estimate `E` separately, ideally from a source that is not you (F-08's independent sources serve).
5. Compute `r = D/P` and `M = rE`.
6. Normalize: `M_norm = ln N * (M − M_min)/(M_max − M_min)` with the range stated in advance.

WORKED NUMBER

D	P	E	r	M	|ln r|	reading
50	50	60	1.000	60.000	0.0000	calibrated contact, M = E
75	50	60	1.500	90.000	0.4055	over-framed by half
30	50	60	0.600	36.000	0.5108	under-framed
50	10	60	5.000	300.000	1.6094	intake collapsing; the P -> 0 limb
50	100	60	0.500	30.000	0.6931	flooding; twice the intake the frame can use

Normalization example: `M = 90` on a declared range `[0, 120]` with `N = 7` gives `M_norm = 1.9459 * 90/120 = 1.4594` nats — 75% of the ceiling.

FAULTS

Fault	Signature	Fix
Ordinal division	`M` computed from star ratings	FVP-2. Demote to the three-way ordering; the numbers already computed are void
Circular `E`	`E` estimated from the same impression that produced `P`	Get `E` from an independent source; otherwise `r` is structurally 1 by construction
Drifting modulus	`M` trends across a block for no reason anyone can name	Freeze the modulus (Protocol step 1) and re-baseline blocks separately
Ceiling violation	`M_norm > ln N`	The normalization range was wrong, or `N` was overstated. Recheck F-03 first

CROSS-REFERENCES. Ch 34 (full treatment), Ch 53 (FVP-2), F-02 (the ratio), F-04 (where `M` goes), F-06 (what `M` funds).

F-02 . THE CALIBRATION RATIO

STATEMENT

```
    r = D / P                          the calibration ratio, dimensionless

    Delta = | ln r |                   the distortion, in nats

   r > 1  OVER-FRAMED   moment exceeds environment  -> projection, seeing the expected
    r = 1   CALIBRATED     moment equals environment   -> contact
       r  <  1     UNDER-FRAMED     moment  falls  short                  -
> flooding, cannot make it mean
```

READS. *How far your frame outruns your intake, measured as a log-ratio so it composes with everything else.*

STATUS. Immediate consequence of F-01; the distortion measure is CANON_INTERNAL §4.

DERIVATION. `r` is the coefficient in `M = rE`. `Delta = |ln r|` is a log-ratio for the same three reasons `Delta_R` is (Chapter 43, §43.0.3): ratios compose multiplicatively, the measure must be invariant to the modulus, and the reading instrument is logarithmic. **`Delta` and `Delta_R` are the same kind of object measured at two places in the circuit — one on the frame, one on the world's answer — and that is why they are directly comparable.**

Independent convergence, stated carefully: predictive-processing accounts model perception as precision-weighting of top-down priors against bottom-up evidence, and **the ratio of prior precision to sensory precision is structurally the same object as `D/P`.** That is a correspondence between formalisms, not an identity claim — predictive processing is an active and contested research program, and structural similarity is not identity.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`r` dimensionless by construction (FVP-1 of F-01 supplies the shared units). `Delta` in nats.
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	Inherits F-01's verdict exactly. **`Delta = |ln r|` is an ordinal quantity pending instrumentation: its sign is a finding, its rank order is a finding, its magnitude is not.** The nats, bits and cents printed below are illustrative constructions showing how the measure would read, not measurements of a person. Ratio scale is restored by Ch 54 Thm 54.3 and by nothing cheaper.

Gate	Verdict	Content
FVP-3 LIMITING CASES	PASS	`r -> inf` (`P -> 0`): total frame dominance. `r -> 0` (`D -> 0`): no frame; `Delta -> inf` in both, correctly, since neither limb has a finite worst case.
FVP-4 FIXED POINTS	PASS	`r = 1`, `Delta = 0`: veridical perception stated as an equation. Under a correction rule `ln r_{t+1} = (1−lambda) ln r_t` (the F-08 dynamics applied to the frame) this fixed point is attracting for `0 < lambda < 2`.
FVP-5 IDENTIFIABILITY	PASS **for `r`**	`r` is identified from `M/E` without separating `D` and `P` — a strictly better position than F-01, and the reason this entry is often more useful than its parent.
FVP-6 FALSIFICATION	SPECIFIED	Predict `M` from `E` and a frozen `r` estimated on a training block; beat the four baselines. **Killed if `r` estimated on block one does not improve prediction on block two over using `r = 1`.**
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do two investigators recover the same `r` for the same person? Unknown.

PROTOCOL

1. Over `n >= 8` encounters, record a predicted magnitude and a realized magnitude on a fixed modulus (this is F-01's protocol run in prediction mode).
2. Compute `ln r_i` for each; take the mean and standard deviation.
3. `r_hat = exp(mean ln r)`. Report the interval, not the point.
4. Compare `Delta = |mean ln r|` against your deadband from F-08 (`ln(1+c)`); a distortion below the deadband is not resolvable.

WORKED NUMBER

Eight illustrative trials — predicted magnitude against realized magnitude on a fixed modulus:

```
predicted: 70  65  80  75  85  60  90  72
realized : 52  50  58  61  60  49  62  55
```

```
mean ln r = 0.2851        sd = 0.0620        SE = 0.0219
r_hat     = 1.3299
95% CI on ln r = [0.2421, 0.3280]  ->  r in [1.274, 1.388]
```

A stable over-frame of about 1.33 with a tight interval that excludes 1 decisively. In the conversion registers: `Delta = 0.2851` nats = 0.4114 bits = 987 cents — a shade under a minor seventh of distortion, and the four-register reading is not decoration; it is a check that the number is in a plausible range for a perceptual quantity.

Reference table:

r	|ln r|	bits	cents	state
0.50	0.6931	1.0000	2400.0	under-framed
1.00	0.0000	0.0000	0.0	**calibrated**
1.25	0.2231	0.3219	772.6	over-framed
1.50	0.4055	0.5850	1403.9	over-framed
2.00	0.6931	1.0000	2400.0	over-framed

Note the symmetry the log gives you for free: `r = 0.5` and `r = 2` are equally distorted. **Under-framing by half and over-framing by double are the same size of error, and only the log-ratio says so.**

FAULTS

Fault	Signature	Fix
Correcting a distortion below the deadband	Effortful frame-work producing no change	Compare `Delta` to `ln(1+c)`. If below, the instrument cannot resolve it; the work is unmeasurable by construction
Direction confusion	Treating flooding as over-framing	`r < 1` is under-framed — too much intake for the frame. The correction is more frame, not less
Single-encounter diagnosis	`r` estimated from one memorable event	`n >= 8` and report the interval (Protocol 3)
Global `r`	One `r` claimed for a whole person	`r` is domain-specific. A person may be over-framed at work and under-framed at home; a pooled `r` averages them into a fiction

CROSS-REFERENCES. Ch 35 (full treatment), Ch 48 (the same ratio at group scope, doctrine-to-data), F-01, F-08.

F-03 . THE INTERNAL SURPLUS

STATEMENT

```
F_int = ln( 1 + (N−1) u )  =  ln N_eff

N       distinguishable categories held by ONE mind
u       non-parallelism between one's OWN framings, in [0,1]
N_eff   = 1 + (N−1)u, the audited count of genuinely different views
F_int   internal surplus, in nats,  0 <= F_int <= ln N
```

READS. *How many genuinely different views you hold, logged — the many windows of the one tower, counted after the audit for redundancy.*

STATUS. Theorem. Paper A, Theorem 2.1, applied within a single observer. Not an analogy: the axioms S1–S4 make no reference to the states belonging to different people.

DERIVATION. `f(u) = ln(1 + (N−1)u)` is uniquely characterized by S1 (geometric dependence), S2 (`f(0) = 0`), S3 (affine effective count), S4 (`f(1) = ln N`). **S3 is the design axiom — it is chosen, not derived, and that must be said wherever the law is leaned on.** Book One's Node 03 HORIZON — three futures cut from three non-parallel glasses — is this theorem applied within one observer, which makes it a derivation rather than an analogy.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`N` a count; u dimensionless in `[0,1]`; `N_eff` a count; `F_int` in nats.
FVP-2 SCALE TYPE	PASS **for u**	`u = 1 − (x.y)^2` is computed from normalized state vectors and is ratio-scale with a true zero (parallel). The **operationalization** of the framings as vectors is where scale trouble enters, and Chapter 39's protocol is the route.
FVP-3 LIMITING CASES	PASS	`u −> 0`: `N_eff −> 1`, `F_int −> 0` — **all your framings are one framing, counted `N` times, at `N` times the cost.** `u −> 1`: `N_eff −> N`, `F_int −> ln N` — orthogonality, the ceiling. `N −> 1`: `F_int = 0` for all u; one category has nothing to be non-parallel to.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	Extremality (Paper A, Thm 4.3): `F = 0` iff parallel, `F = ln N` iff orthogonal. Both are attained, both are meaningful, and the lower one is the collapse this manual spends Chapters 11 and 22 preventing.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`F_int` is identified given `N` and `u`. `N` and `u` **trade off**: a small `N` at high `u` and a large `N` at low `u` produce identical `F_int`. Separating them requires counting the categories independently of measuring their overlap.
FVP-6 FALSIFICATION	SPECIFIED	Paper A is pure mathematics and makes no empirical claim; what is falsifiable is the **operationalization**. Preregister the framing-elicitation procedure, fit, freeze, and test whether measured `u` predicts held-out performance on a task requiring multiple framings, against the four baselines including shuffled-input.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do two investigators eliciting one person's framings recover the same `u`? This is exactly the specific-objectivity question of Chapter 54 and it is open.

PROTOCOL

1. Elicit framings **independently and in writing before any are compared** — the personal-scope form of Chapter 17's silent-write-first rule, which prevents the framings from converging during elicitation.
2. Count them: that is `N`, and it is an upper bound, not a measurement.
3. Estimate pairwise overlap; `u` is one minus the squared cosine, averaged or taken pairwise per Chapter 39.
4. Compute `N_eff` and `F_int`. **Report `N_eff`, never `N`, when describing how many views you hold.**

WORKED NUMBER

N = 7 framings, ln N = 1.9459 nats:

u	N_eff	F_int	% of ceiling	marginal dF/du
0.00	1.000	0.0000	0.0%	6.0000
0.10	1.600	0.4700	24.2%	3.7500
0.30	2.800	1.0296	52.9%	2.1429
0.45	3.700	1.3083	67.2%	1.6216
0.60	4.600	1.5261	78.4%	1.3043
1.00	7.000	1.9459	100.0%	0.8571

The marginal value of the k-th genuinely orthogonal framing is ln k − ln(k−1):

```
k=2  0.6931 nats (1.0000 bits)     k=7   0.1542 (0.2224)
k=3  0.4055        (0.5850)        k=8   0.1335 (0.1926)
k=4  0.2877        (0.4150)        k=9   0.1178 (0.1699)
k=5  0.2231        (0.3219)        k=10  0.1054 (0.1520)
k=6  0.1823        (0.2630)        k=12  0.0870 (0.1255)
```

The second genuinely different framing is worth one full bit — the largest single gain you will ever buy. The twelfth is worth 0.126 bits, an eighth as much, and Chapter 39 calls the fortieth glazing bills. Both facts are the same logarithm, and neither is pessimism: concavity is the lens (CANON_INTERNAL §2).

FAULTS

Fault	Signature	Fix
Counting N instead of N_eff	"I look at this from six angles"	Measure u. Six angles at u = 0.10 is N_eff = 1.5
Pseudo-diversity	Framings differ in vocabulary, agree in prediction	Run the collision test: do the framings ever disagree about an outcome? If not, u ~ 0
The N/u trade	Claiming a high F_int without stating which term produced it	FVP-5. Report both, always
Ceiling denial	Effort spent raising u when u is already 0.9	At u = 0.9, F_int is 95.4% of ceiling; the remaining move is to raise N (Ch 8)

CROSS-REFERENCES. Ch 39 (full treatment), Ch 7, Ch 8, Ch 11, Ch 22; F-01 (ln N bounds M), F-05 (h is eta times the unclaimed part of this).

F-04 . THE SELF-STORAGE LAW

STATEMENT

```
Q_{t+1} = (1 − delta) Q_t + eta ∗ M_t − C_t
```

```
   Q_t    stored order: character, skill, memory, internalized structure  [Q-
units]
   M_t    the moment as lived, normalized to [0, ln N]                 [nats]
   delta  retention loss per period, in [0,1]                       [1/period]
    eta    conversion of experience into durable structure, >= 0         [Q-
units/nat]
    C_t    cognitive load, effort, maintenance of self, >= 0             [Q-
units/period]

   Admissibility:     C_t <= (1-delta) Q_t + eta M_t   guarantees Q_{t+1} >= 0
    One-step growth:    Q_{t+1} > Q_t  <=>  eta M_t > delta Q_t + C_t
```

READS. *Character is what is left of every moment you have had, after the leak, the conversion loss, and the cost of maintaining yourself.*

STATUS. Theorem. Paper E §3, in the internal-scope form CANON_INTERNAL §8 specifies. **The join between the two formulas: the Moment feeds the storage law as its input.**

DERIVATION. Direct substitution of `M_t` for `S_t`, licensed by CANON_INTERNAL §1's forced consequence that `M` is bounded by `ln N` and so shares `S_t`'s range. **The scale-free claim — that a mind and a city run the same equation with different parameters — is a claim about MODEL FORM, badged as such in CANON §1.4.**

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	Every term in Q-units per period once `eta` carries Q-units/nat. `delta Q_t` is Q-units/period. Both sides agree. **`eta`'s units are the conversion key and stating them is what stops the equation from being a slogan.**

Gate	Verdict	Content
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	`Q` must be at least interval-scale for the difference terms and ratio-scale for `Q_inf` percentages. **Badged: that `Q` is a scalar at all — rather than a vector of incommensurable stores — is a postulate of this manual (Appendix E), not a finding.** And the input term inherits F-01's demotion: **the `M -> Q` join is structural, not yet quantitative.** The recursion, the two-track ledger, the fixed point and the orderings stand; numbers obtained by feeding a measured `r` through `M` do not. Rasch calibration (Ch 54 Thm 54.3, ~123 items per term) is the route to satisfying both.
FVP-3 LIMITING CASES	PASS	`delta -> 0`: no fixed point; `Q_t <= Q_0 + t eta ln N` (Thm 3.3) — the tradition that never forgets and never settles. `delta -> 1`: `Q_{t+1} = eta M_t - C_t`, total amnesia, the state is only the present. `eta -> 0`: `Q` decays to `-C/delta` — **rumination is recursion with `eta ~ 0`.** `C -> 0`: no maintenance cost, `Q_inf = eta M/delta`.
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* - c*)/delta` (Thm 3.5). **If `eta M* < c*` the fixed point is negative and the regime is unsustainable or the linear model must be replaced. The theorem does not hide this and neither does this entry.**

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL**	From {Q_t} alone, delta is identified from successive gaps and the composite eta M∗ − c∗ from delta Q_inf. **eta, M∗, c∗ are not separately identified.** Measure two of the three independently (F-01 for M, Ch 5's ledger for C) and recover the third.
FVP-6 FALSIFICATION	SPECIFIED	Paper E §5.3. Preregister period, Q instrument, window. Fit delta and eta on training, **freeze**, forecast held-out Q. Beat persistence, direct-input, matched AR(1), shuffled-input control. **Killed if the frozen model does not beat the shuffled-input control**, since that control destroys exactly the accumulation the law claims.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do delta and eta recover across people and across domains within a person? Unknown.

PROTOCOL

1. Declare the **period** and hold it.
2. Instrument Q — a calibrated measure of the stored capacity in question, not a mood rating.
3. Log M_t per F-01 and C_t per Chapter 5's cost ledger, in the same units.
4. Fit delta from the decay of Q across a deliberate zero-input window (a period with M ~ 0), which is the cleanest identification available: with M = 0, Q_{t+1} = (1−delta)Q_t − C_t.
5. Recover eta by residual, and **report it with the identifiability caveat attached**.
6. Check admissibility every period. A violated admissibility condition is a broken model, not a bad week.

WORKED NUMBER

Illustrative: N = 7, u = 0.45 (so M∗ = F_int = 1.3083), eta = 0.30, delta = 0.10, C = 0.25, Q_0 = 0.80.

```
    Q_inf = (0.30 * 1.3083 - 0.25) / 0.10 = 1.4250 Q-units
   one-step growth test at Q = 1.5:  eta*M = 0.3925  vs  delta*Q + C = 0.4000  -
>  NO GROWTH
```

```
sustainability ceiling at Q = 1.5 (Cor 3.4):  C <= eta lnN − delta Q = 0.4338
```

A ten-period trajectory with a varying moment, including two poor periods:

t	M_t	Q_t	step dQ
0	0.45	0.8000	-0.1950
1	1.10	0.6050	+0.0195
3	0.20	0.7771	-0.2677
4	1.60	0.5093	+0.1791
5	1.60	0.6884	+0.1612
6	1.60	0.8496	+0.1450
7	0.90	0.9946	-0.0795
9	1.60	1.0536	+0.1246
10	—	1.1783	—

One bad period at `t = 3` (`M = 0.20`) cost 0.2677 Q-units, which took the next two good periods to recover. That asymmetry is not psychological; it is `eta M − delta Q − C` evaluated at a low `M`, and it is why cadence beats intensity: the leak runs whether or not the input does.

FAULTS

Fault	Signature	Fix
`eta ~ 0` recursion (rumination)	High `M`, high effort, `Q` flat or falling	Nothing is converting to structure. Change the *form* of the practice (writing, artifact, place), which is what `eta` is
`delta ~ 0` fantasy	A plan in which nothing decays	Name the loss mechanism or the model does not apply (Ch 44, CS-3)
Negative fixed point run anyway	`eta M* < c*`, `Q` declining, practice continued	Compute `Q_inf` before starting. Cut `C`, raise `eta` or `u`, or stop
Mood as `Q`	`Q` instrument is a daily feeling rating	Mood is not stored order. `Q` is what remains when the mood has gone
Ignoring admissibility	Model produces `Q < 0` and is still trusted	A violated admissibility condition invalidates the step, not the week

CROSS-REFERENCES. Ch 40 (full treatment), Ch 2 (the leaky integrator), Ch 3 (`delta`), Ch 4 (`eta`), Ch 5 (`C`); F-01 (the input), F-05, F-09.

F-05 . PERSONAL HEADROOM

STATEMENT

```
h_t = eta * ln N - delta * Q_t - C_t                [Q-units/period]

h_t < 0   unsustainable in this model. No geometry saves it.
h_t = 0   maintenance only. Hold, do not build.
h_t > 0   growth available; any requested increment r_t must satisfy r_t <= h_t.

At any fixed point:   h_inf = eta * ( ln N - F(u*) )    THE ARRIVAL INVARIANT
```

READS. *How much capacity for the new you have this period, if your geometry were perfect.*

STATUS. Theorem. Paper E §6.4 and Corollary 4.2, at internal scope. The arrival invariant is derived in Chapter 44, §44.4.

DERIVATION. h is the growth condition `eta M_t > delta Q_t + C_t` evaluated at the **maximum possible** moment `M = ln N`. It is therefore the margin at perfect geometry. **Consequence that must travel with the entry: `h > 0` is necessary and not sufficient for growth, because the actual step uses `F(u*)`, not `ln N`.** The shortfall — call it the geometry gap — is `eta(ln N - M_t)`, and it is exactly the quantity the arrival invariant names.

THE CONCAVITY CORRECTION, and it applies to every feasibility check made with this entry. The protocol below asks whether a requested increment `r_t` fits inside `h_t`. That test is run against a surplus figure, and the surplus figure must be the **mean of `F(U_t)`, never `F` of the mean `U`.** What accumulates over periods is `sum of eta F(U_t)`, and because `F` is strictly concave (Paper A, Theorem 3.3), Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )      with equality only if U is constant
```

A system whose geometry varies delivers strictly less surplus than one that sits at its own average every period, and testing feasibility on `F(mean U)` therefore approves increments that cannot be delivered. The gap is the concavity penalty. In the worked office of Chapter 26 §26.7 it is 0.0761 nats — mean `F` of 1.1501 against `F(mean U)` of 1.2262 — **a loss of 6.2 per cent of the surplus purely to variability**, and it was large enough there to turn a feasible-looking plan into a shortfall. Report `U_req` for legibility if you wish; do not test against it. Regularising a variable geometry recovers the penalty for free, which makes it the cheapest surplus available to most operators.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	All three terms in Q-units/period. `eta ln N` is Q-units/nat times nats.
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	Inherits F-04's requirement on `Q` and Ch 5's requirement that `C` be commensurated into Q-units — **which is itself a postulate of this manual, badged in Appendix E: that costs of different kinds can be put on one scale.**

Gate	Verdict	Content
FVP-3 LIMITING CASES	PASS	`Q -> 0`: `h -> eta ln N − C`, the maximum headroom a person can have. `Q -> (eta lnN − C)/delta`: `h -> 0`, the ceiling on stored order beyond which maintenance consumes everything. `C -> eta ln N`: `h <= 0` for any `Q > 0`.
FVP-4 FIXED POINTS	PASS	`h = 0` at `Q = (eta ln N − C)/delta` — the **maintenance ceiling**, distinct from and above `Q_inf`. And the arrival invariant: at `Q_inf`, `h = eta(ln N − F(u∗))`, independent of `delta` and `C`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	Inherits F-04's non-separation of `eta`, `M∗`, `c∗`. `h` itself is computable once `eta`, `N`, `delta`, `Q`, `C` are in hand; the burden is on `N` and on the commensuration of `C`.
FVP-6 FALSIFICATION	SPECIFIED	Preregister; predict whether a requested increment `r_t` is achieved, from `h_t` computed in advance. **Killed if achieved increments routinely exceed `h_t`** — that outcome would falsify Corollary 4.2's claim that no policy based on interaction geometry alone can guarantee more.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Cross-individual recovery unrun.

PROTOCOL

1. Obtain `N` (F-03), `eta` and `delta` (F-04), `Q_t` (F-04's instrument), `C_t` (Chapter 5).
2. Compute `h_t` each period and plot it. **`h` is the gauge; `Q` is the tank.**
3. Before accepting any new commitment, express its expected increment `r_t` in Q-units and check `r_t <= h_t`.
4. At a settled state, compute `h_inf = eta(ln N − F(u∗))` and account for the remainder.

WORKED NUMBER

`eta = 0.30`, `N = 7` (`ln N = 1.9459`), `delta = 0.10`, `C = 0.25`:

Q	h	verdict
0.00	0.3338	growth available
1.00	0.2338	growth available
2.00	0.1338	growth available
3.00	0.0338	growth available, barely
3.34	0.0000	**maintenance ceiling**
4.00	-0.0662	unsustainable

`h = 0` at `Q = (0.30 * 1.9459 − 0.25)/0.10 = 3.3377`. At the fixed point of the F-04 example (`u = 0.45`, `Q_inf = 1.4250`):

```
h_inf = eta ( ln N − F )  = 0.30 * (1.9459 − 1.3083) = 0.1913
check from the gauge      = 0.30*1.9459 − 0.10*1.4250 − 0.25 = 0.191273
```

Identical to six decimal places, and it does not contain `delta` or `C`. The reading: at arrival, this operator has 0.1913 Q-units per period of unclaimed capacity, and it is unclaimed for exactly one reason — their framings claim only 67.2% of the ceiling. **Large headroom at a settled state is not a reserve. It is a report that your framings are parallel.**

FAULTS

Fault	Signature	Fix
Headroom as reserve	`h > 0` read as "I have capacity to spare"	`h_inf = eta(ln N − F)`. Raise `u` or `N`; do not "spend the slack"
`h > 0` read as sufficient	Growth expected because `h > 0`, but `Q` falls	`h` is the margin at `M = ln N`. Actual growth needs `eta M_t > delta Q_t + C_t`
Feasibility tested on `F(mean U)`	A plan that clears the gate on paper and falls short every period	Test on **mean `F(U_t)`**. `F` is strictly concave, so `F` of the mean overstates; the gap was 6.2% in the Ch 26 case
Uncommensurated `C`	Only monetary or only time costs counted	Chapter 5's five ledgers. An omitted cost inflates `h` and licenses a commitment that cannot be met
Negative `h` treated as motivational	"I need to push harder"	Cor 3.4: **no choice of interaction geometry prevents order from decreasing.** Cut `C`, shed `Q`, or raise `eta`

CROSS-REFERENCES. Ch 6 (full treatment at external scope), Ch 40, Ch 5; F-03, F-04, F-09.

F-06 . THE LAUNCH BUDGET

STATEMENT

```
L = max( 0, M − C )                    [nats, or Q-units once eta is applied]

L    the surplus available for imagining, planning, or attempting elsewhere
M      the moment as lived, in [0, ln N]
C      the cost of the period, commensurated into the same units

L = 0            no surplus to dream with
L_max = ln N − C the ceiling, when the moment is maximal
```

READS. *Escape is dreamt only out of surplus; when the moment costs everything it produces, there is nothing left to imagine elsewhere with.*

STATUS. Derived formula of this manual (CANON_INTERNAL §8), formalizing Book One's Node 00 departure-lounge doctrine. Derived, with one badge.

DERIVATION. The storage law's one-step content is `eta M_t − C_t`: what a period contributes before the leak takes its share. `L` is that quantity's positive part, before conversion — the *unallocated* portion of a moment, the part not consumed by the cost of having it.

Postulate L1 (the launch condition). Imaginative and exploratory activity draws on the same budget as maintenance, and draws second. **Badged: an assumption about the ordering of claims on a period's surplus, not a result.** It is what makes `L` a max-with-zero rather than a quantity that can go negative and be carried, and it is consistent with the departure-lounge doctrine — *in those minutes there was no surplus to dream with* is `L = 0` — but the ordering is unmeasured.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M` in nats; `C` must be converted into nats (divide the Q-unit cost by `eta`) or `M` converted into Q-units (multiply by `eta`). **The conversion must be stated; mixing the two is the commonest arithmetic error in this entry.**

Gate	Verdict	Content
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	Subtraction needs interval scale on a common unit; the commensuration of `C` is again a badged postulate (Appendix E). **`M` additionally inherits F-01's demotion, so `L` is reportable as a threshold condition — is the moment above the cost line or not — and not as a magnitude.** The sign of `M − C` is the finding; its size is not, until Ch 54 Thm 54.3's instrument exists.
FVP-3 LIMITING CASES	PASS	`M <= C`: `L = 0` — the departure lounge with nothing to depart with. `M −> ln N`: `L −> ln N − C`, the ceiling. `C −> 0`: `L −> M`; the whole moment is available. `C −> ln N`: `L = 0` for every possible moment — **a cost equal to the ceiling makes the person structurally incapable of imagining elsewhere, whatever they experience.**
FVP-4 FIXED POINTS	PASS	The boundary `M = C` is the **no-launch surface**. `L` is continuous but **not differentiable** there — a genuine kink, and the non-smoothness is the phenomenon: the transition from "some budget" to "none" is abrupt, not gradual, which is why people describe it as a wall rather than a slope.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`L` is identified from the difference; `M` and `C` separately are not, if only their difference is observed. Since the intervention differs completely (raise `M` vs cut `C`), **the separation is required before acting, not before measuring.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Preregister `M` and `C` measurement; predict, from a frozen model, which periods produce exploratory output (plans, applications, drafts, attempts). Beat the four baselines. **Killed if exploratory output is uncorrelated with `L` after controlling for `M` alone** — that would say the cost side carries nothing and the doctrine is wrong.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Unrun.

PROTOCOL

1. Measure `M_t` per F-01, normalized to `[0, ln N]`.
2. Measure `C_t` per Chapter 5's ledger and convert to nats by dividing by `eta`.
3. `L_t = max(0, M_t − C_t)`. Log it every period.
4. Report **the count of `L = 0` periods**, not just the mean. The zeros are the finding.

WORKED NUMBER

M	C	L	state
0.20	0.60	0.000	no surplus to dream with
0.45	0.45	0.000	exactly at the surface
0.60	0.45	0.150	launch available
1.90	0.45	1.450	launch available
1.90	1.95	0.000	high moment, higher cost — **the busy trap**

The last row is the entry's most important case. **A person can have an extremely rich moment and zero launch budget, and from the outside they look like someone who has everything and does nothing with it.** The arithmetic says the opposite: they have nothing left, because `C` exceeded `M`.

Over the F-04 example's ten periods with `C = 0.25`:

```
sum L = 9.75      mean L = 0.975      periods with L = 0: 1
L_max = ln N − C  = 1.9459 − 0.25 = 1.6959
```

FAULTS

Fault	Signature	Fix
Order-of-operations reversal	"Get out of your comfort zone" prescribed to someone at `L = 0`	The dream is manufactured *from* surplus. Raise `M` or cut `C` first; the exhortation has no budget to act on (Ch 41)
The busy trap	High `M`, high `C`, `L = 0`, and self-blame for lack of ambition	Cut `C`. The moment is fine; the ledger is not
Unit mixing	`M` in nats minus `C` in Q-units	Convert. Divide the Q-unit cost by `eta`
Mean without zeros	Reporting mean `L` across a block containing many zeros	Report the zero count. A mean of 0.5 from alternating 1.0 and 0.0 is a different life from a steady 0.5

CROSS-REFERENCES. Ch 41 (full treatment), Ch 5, Ch 42; F-01, F-05.

F-07 . THE GRADIENT

STATEMENT

```
    g = − grad C                              the desire line, a vector field

   C(x)   the local cost of being at / doing configuration x  [Q−units/period]
  x    a configuration coordinate: actions−to−begin, distance, latency, friction
   g      points in the direction of cheapest motion

   Observable proxy, under Postulate G1:
       n(x) ~ exp( −C(x) / theta )   ->   ln n(x) = const − C(x)/theta
       fitted slope b = −(1/theta) dC/dx
```

READS. *Water does not negotiate; read the worn path rather than arguing with it.*

STATUS. Derived formula of this manual — Book One's Node 02 GRADIENT formalized. Derived, with one badge.

DERIVATION. The storage law's cost term `C_t` is a scalar per period, but the cost of a period depends on what was done in it; treat `C` as a function over a configuration space and the direction of cheapest motion is `−grad C`. That much is definitional. **The empirical content requires a link between cost and behaviour, and that link is a postulate:**

Postulate G1 (the choice rule). The frequency with which a configuration is entered is log-linear in its local cost: `n(x) ~ exp(−C(x)/theta)`, `theta > 0` a choice temperature. **Badged: the standard log-linear choice form, adopted as a modelling assumption — not a result of this corpus and not a claim about any measured population.** It puts behaviour in the same logarithmic register as everything else in the book, at the price of a free `theta` (see FVP-5).

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`C` in Q-units/period; `x` in whatever the coordinate is (actions, metres, seconds); `grad C` in Q-units per period per unit of `x`. Under G1 the fitted `b` is in `(1/theta)` times those units — **so the fitted slope is a gradient in units of the choice temperature, not in Q-units.**
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	`C` must be ratio-scale (Ch 5). `x` must be at least interval-scale; counts of actions are ratio-scale and are therefore the preferred coordinate. `n` is a count — ratio-scale — which is what makes the log-linear fit legitimate.
FVP-3 LIMITING CASES	PASS	`grad C -> 0`: no desire line; motion is undirected — the flat that Book One's Node 00 calls the stillpoint. `|grad C| -> inf`: an infinitely steep wall; the configuration is never entered. `theta -> 0`: choice becomes deterministic, always the cheapest option. `theta -> inf`: choice becomes uniform, cost carries nothing.
FVP-4 FIXED POINTS	PASS	`grad C = 0` at local minima: **the comfort zone, defined without sentiment as a stationary point of the cost field.** Book One's Node 00 STILLPOINT — *you are already there* — is the statement that the present configuration is at or near such a point by construction.

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL — and instructively so**	From behaviour alone, only `C/theta` is identified, never `C` and `theta` separately. **But the *gradient* is better identified than the *level*: adding any constant to `C` changes no behaviour and no fitted slope.** So this entry can report reliable differences and cannot report absolute costs. Say so whenever a fitted `b` is quoted.
FVP-6 FALSIFICATION	SPECIFIED	Preregister the coordinate and the count window. Fit `b` on a training block, **freeze**, then **intervene**: remove `k` actions from one configuration and predict the new count as `n * exp(-b k)`. Beat persistence, direct-input, matched AR(1), shuffled-input. **Killed if the predicted response to a manipulated coordinate is not observed** — this is the one entry in the formulary with a cheap, direct, interventional test, and it should be run first.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `b` recover across people and across domains? Unknown.

PROTOCOL

1. Choose a coordinate that is a **count**: number of discrete actions required to begin.
2. For `>= 5` activities, record the action count `k_i` and the initiation count `n_i` over a fixed window.
3. Fit `ln n = a + b k` by least squares. Report `b`, `R^2`, and the fitted values.
4. Compute the **halving distance**: `ln(0.5)/b` actions — the number of added actions that halves initiation.
5. **Intervene and re-measure.** Step 5 is not optional; it is FVP-6.

WORKED NUMBER

Illustrative log over one week — six activities, actions-to-begin against initiation count:

actions k	observed n	fitted n
1	26	25.05
2	17	16.62
3	11	11.03
6	3	3.22
9	1	0.94

```
fit:  ln n = 3.6308 − 0.4100 k        R^2 = 0.9982
gradient magnitude |b| = 0.4100 nats of initiation per added action
halving distance = ln(0.5)/(−0.4100) = 1.69 actions
removing 2 actions multiplies initiation by exp(2∗0.4100) = 2.271
```

Every 1.69 actions added to the front of a practice halves how often it happens. That is the whole of Book One's second node in one number, and it is the number that makes threshold design (Chapter 15) an engineering discipline rather than an aesthetic one.

The honest caveat, from FVP-5: **0.4100 is a gradient in units of the choice temperature, and the absolute cost of an action is not recoverable from this fit.** What is recoverable — and what the intervention test checks — is that removing two actions should multiply initiation by about 2.27.

FAULTS

Fault	Signature	Fix
Arguing with the gradient	Exhortation applied to a configuration nobody enters	Regrade. Remove actions. The path is data, not an opinion to be corrected
Level claims from slope data	"This costs me 0.41 Q-units"	FVP-5. Only differences are identified. Report ratios of frequencies, never absolute costs
Ordinal coordinate	x is "difficulty, 1-5"	Use a count. Actions, seconds, metres
No intervention	Fit reported, never tested	FVP-6. The manipulation is cheap and it is the only real test this entry has
Confounded coordinate	High-k activities also happen to be the disliked ones	Manipulate k on **one** activity and hold the activity fixed; that is what makes it an experiment

CROSS-REFERENCES. Ch 42 (full treatment), Ch 15 (thresholds), Ch 26 (intervention); F-05, F-06.

F-08 . THE RETURN DIFFERENCE

STATEMENT

```
rho_t     = ln( M'_t / M_t )                signed return, in nats
Delta_R,t = | rho_t |                       magnitude, in nats

M_t   the moment sent       M'_t  the moment returned

rho > 0  OVER-RETURN   |  rho = 0  NULL RETURN  |  rho < 0  UNDER-RETURN

deadband:    Delta_R_min = ln(1 + c)       c = the operator's Weber fraction
correction:   rho_{t+1} = (1 - lambda) rho_t  lambda = correction gain
stability:    convergent iff 0 < lambda < 2
```

READS. *The world's answer is never the voice you sent, and the difference is the only data in the exchange.*

STATUS. Derived formula of this manual (CANON_INTERNAL §8), Book One's Node 08 ECHO. Derived, with one badge (Postulate R1, the returnable moment — Chapter 43, §43.0.2).

DERIVATION. Chapter 43 in full. In one line: the corpus works in log-ratios everywhere, so the only difference measure that composes with the rest of the formulary is the log-ratio, and it is forced rather than chosen.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M'/M` dimensionless; `rho` in nats, convertible by the §45.0.4 keys.
FVP-2 SCALE TYPE	**DEMOTE — binding** (Standing Verdict 1)	Ratio scale with a true zero, or the entry reads as a **sign-only ordering**. `M` and `M'` are moments and inherit F-01's demotion, so **that is where the entry stands today: over-return, null return and under-return are findings; the number of nats is not.** A `rho` computed from Likert data was always void rather than imprecise; a `rho` computed from an uncalibrated magnitude estimate is merely unreportable. Ch 54 Thm 54.3 supplies the missing scale, at ~123 items per term for half a logit.

Gate	Verdict	Content
FVP-3 LIMITING CASES	PASS	`M' -> M`: `rho -> 0`, no information (contact or closure). `M' -> 0`: `rho -> -inf`, total non-return — being ignored, and the divergence is correct because there is no finite penalty scale for it. `M -> 0`: undefined, correctly — an echo of nothing is not an echo.
FVP-4 FIXED POINTS	PASS	`rho = 0`, unique and attracting under the correction map for `0 < lambda < 2`. Meaning: agreement between your construction and the world's answer — **which is either contact or closure, and only the scatter test distinguishes them.**
FVP-5 IDENTIFIABILITY	**PARTIAL**	`rho` recoverable from paired estimates. The **decomposition** into frame error and world response is **not identified from `rho` alone**; separating them requires an independent estimate of `M'` or a deliberate manipulation of `D`.
FVP-6 FALSIFICATION	SPECIFIED	Fit `lambda` on training, freeze, predict held-out `rho`. Beat persistence, direct-input, matched AR(1), and the **shuffled-input control** (same `M` series, permuted `M'`). Killed by failure against the shuffled control.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do `lambda` values recover across people? Unknown; Chapter 43 §43.11 states the design.

PROTOCOL

1. Fix the modulus; record `M` **before** the return arrives (the single most violated rule).
2. Record `M'` from `>= 2` sources asked **independently**, timestamped to the exchange they answer, not to their arrival.

3. Estimate `sigma_rho` from `>= 8` equivalent exchanges before interpreting anything.
4. Declare the deadband `ln(1+c)` and refuse to read below it.
5. Compute `rho`, `Delta_R`, and per block: `R_net = sum rho`, `R_gross = sum |rho|`, `E_R = |R_net| / R_gross`.
6. After four consecutive returns inside the deadband, run the flattery test.
7. **An echo source may decline; a decline is missing data, never a null return.**

WORKED NUMBER

M	M'	ratio	rho	Delta_R	bits
60	60	1.0000	0.0000	0.0000	0.0000
60	63	1.0500	+0.0488	0.0488	0.0704
60	45	0.7500	-0.2877	0.2877	0.4150
60	90	1.5000	+0.4055	0.4055	0.5850
60	20	0.3333	-1.0986	1.0986	1.5850

Deadband, at three Weber fractions: `c = 0.05` → `0.0488` nats; `c = 0.10` → `0.0953`; `c = 0.20` → `0.1823`. **An operator with `c = 0.20` cannot distinguish the second row of that table from silence.**

Detection cost, `n >= (1.96 sigma/Delta)^2`: at `sigma_rho = 0.30`, resolving a `Delta_R` of 0.05 takes **139 echoes**; resolving 0.20 takes **9**.

FAULTS

Fault	Signature	Fix
Zero read as success	`Delta_R ~ 0` celebrated	Run the scatter test. Converged contact varies; closure does not
The flattering mirror	Four-plus returns inside the deadband, no sign changes, all sources agree	Add a genuinely independent source; compute `n_eff`
The devouring echo	`|rho|` growing, signs alternating with rising amplitude	`lambda >= 2`. Reduce below 1 and hold
Post-hoc `M`	`rho` implausibly near zero across a whole log	Record `M` before the return. Compare a pre-recorded block against a reconstructed one
Worth conversion	Negative `rho` narrated as a statement about the self	FVP-5: the decomposition is not identified. The instrument does not license the claim
Correcting faster than the echo returns	Oscillation appearing only when cadence increased	Correction cadence <= echo cadence

CROSS-REFERENCES. Ch 43 (full treatment), Ch 11, Ch 18, Ch 25; F-01, F-02, F-09.

F-09 . THE CONTRACTION AND ARRIVAL

STATEMENT

```
| Q_t − Q_inf | = (1 − delta)^t | Q_0 − Q_inf |            THE ERROR LAW

Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )            the whole trajectory

 Q_inf = ( eta ∗ M∗ − c∗ ) / delta                          Paper E, Thm 3.5
 tau   = −1 / ln(1 − delta)                          periods per factor of e
t(eps) = ln(eps) / ln(1 − delta)                    periods to close to a fraction eps

a−posteriori distance:  | Q_t − Q_inf | = ((1−delta)/delta) ∗ (last step into Q_t)
```

READS. *Where you are is where you are going, plus what is left of where you started — and what is left of where you started is halving on a fixed schedule.*

STATUS. Theorem. Paper E Theorem 3.5 plus the error law derived from Theorem 3.1; Banach's fixed point theorem supplies existence, uniqueness, and robustness. **Book One's Node 10 FIXEDPOINT is Theorem 3.5, and Chapter 44 proves the correspondence rather than asserting it.**

DERIVATION. Chapter 44, §44.0.2–§44.0.5. `T(Q) = (1−delta)Q + b` is a contraction on the complete space `(R, |.|)` with constant exactly `1−delta` whenever `delta > 0`; Banach gives the unique fixed point, and subtracting the fixed-point equation from the update gives the error law by induction. **The leak is what makes arrival possible: at `delta = 0` the map is an isometry, there is no fixed point, and `Q_t = Q_0 + tb` grows without bound.**

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`Q` in Q-units; `(1−delta)^t` dimensionless; `tau` in periods. **The period must be declared or none of it means anything.**
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	Interval scale for the differences, ratio scale for percentage readings. Inherits F-04's badged postulate that `Q` is a scalar.

Gate	Verdict	Content
FVP-3 LIMITING CASES	PASS	`t -> inf`: gap -> `0`. `delta -> 0+`: no convergence, `Q_inf -> inf` unless `b = 0` — the system that never forgets and never settles. `delta -> 1`: arrival in one step; the state is only the present drive. `Q_0 = Q_inf`: gap zero for all `t` — you were already there.
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* - c*)/delta`, unique by Banach, **containing no `Q_0`** — which is why arrival is recognition rather than victory. The error map `e -> (1-delta)e` has the unique fixed point `e = 0`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`delta` identified from successive gaps; the composite `eta M* - c*` from `delta Q_inf`; the three components **not separately identified** without independent measurement of two of them.
FVP-6 FALSIFICATION	SPECIFIED	Fit `delta` on training, freeze, forecast held-out `Q`. Beat persistence, direct-input, matched AR(1), shuffled-input. **Killed if the frozen geometric-decay model does not beat the shuffled-input control**, since that control destroys the ordering of inputs, which is the entire content of "it accumulates".
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `delta` recover across people, investigators, domains? Unknown; Chapter 44 §44.10 states the design.

PROTOCOL

1. Declare the period.

2. Compute `Q_inf` **before** pursuing it, from `eta`, `M*`, `c*`, `delta`.
3. Check the sign. `eta M* < c*` gives a negative fixed point: the loop converges toward destruction, and this must be reported before the practice is run, not discovered during it.
4. Declare the arrival criterion `c` and compute `t_arrive = ln(c)/ln(1−delta)` from a standing start.
5. Track distance from arrival each period as `((1−delta)/delta)` times the last step.
6. At arrival, compute `h_inf = eta(ln N − F(u*))` (F-05) and account for it.

WORKED NUMBER

`delta = 0.10`, `Q_inf = 4.0`, `Q_0 = 0.5`:

t	gap	Q_t
0	3.50000	0.50000
22	0.34467	3.65533
29	0.16485	3.83515
44	0.03394	3.96606
66	0.00334	3.99666

```
to 5%: 29 periods      to 1%: 44 periods      to 0.1%: 66 periods
tau = 9.49 periods     half-life = 6.58 periods
```

Each halving costs 6.58 periods and each factor of ten costs 21.85 periods, and both are constant for the entire trajectory. Closing the gap from 90% to 99% costs 22 periods; closing it from 99% to 99.9% also costs 22 periods; closing the first half costs 7. **That is the arithmetic behind every practice that felt fast at the start and interminable at the end, and nothing slowed down.**

The full convergence table across `delta` is Chapter 44, §44.1; the four landmarks are `delta = 0.05 -> 59 periods to 5%`, `0.10 -> 29`, `0.20 -> 14`, `0.30 -> 9`.

FAULTS

Fault	Signature	Fix
The unarrived asymptote	"Almost there" for many periods, read as failure	There is no arrival event. Declare `c` in advance
Effort substitution	More effort, no change in `eta`, `u`, `C`, `delta`; same path	Effort acts on the transient. Name which parameter the effort moves; if none, stop
Cadence confusion	Cadence increased, calendar unchanged	`tau` depends on per-period `delta`. Cadence acts through `eta` and `C`
Counterfeit arrival	`Q` rising with no change in your own four parameters	The extra drive is a transfer from someone. Run the audit on every party (Ch 44, §44.7.4)
Stall named arrival	Flat `Q`, no `Q_inf` on record, no review date	The arithmetic cannot adjudicate; the documentation can

Fault	Signature	Fix
Low-`delta` noise	Very stable practice, large swings traced to small perturbations months earlier	`1/delta` amplification. Accept it or shed deliberately

CROSS-REFERENCES. Ch 44 (full treatment and the proof), Ch 1, Ch 2, Ch 23; F-04, F-05, Appendix F.

The Specification

How to run the formulary. **These are ordering and hygiene requirements; the per-entry protocols sit inside them.**

PF-1 — Run in dependency order. `F-03 -> F-01 -> F-02 -> F-08 -> F-04 -> F-05 -> F-06 -> F-07 -> F-09`. `N` and `u` first because they set the ceiling everything else lives under; the contraction last because it needs every other parameter.

PF-2 — One modulus, one period, one block. All magnitude estimates in a block share a modulus; all rates share a period; both are declared in writing before data collection and are not changed mid-block.

PF-3 — Zero-input window first. Fit `delta` during a deliberate low-input period before fitting anything else. It is the only clean identification of `delta` available at internal scope, and everything downstream depends on it.

PF-4 — Pre-record, never reconstruct. `M` before the return (F-08); `Q_inf` before the pursuit (F-09); the arrival criterion before arrival. **A formulary run on reconstructed inputs produces internally consistent numbers that are jointly wrong, which is the worst available outcome.**

PF-5 — Report intervals, not points. Every estimate carries an `n` and a spread. A point estimate with no `n` is not a measurement.

PF-6 — Compare each quantity to its deadband. `Delta` (F-02) and `Delta_R` (F-08) against `ln(1+c)`; `dQ` (F-04) against the `Q` instrument's resolution. **Below the deadband, report "unresolved", never "zero".**

PF-7 — Carry the identifiability caveats into every conclusion. The non-identifications are: `D` versus `P` (F-01); `N` versus `u` (F-03); `eta`, `M*`, `c*` (F-04, F-05, F-09); frame error versus world response (F-08); `C` level versus `theta` (F-07). **A conclusion requiring one of these separations must name the second measurement that supplied it.**

PF-8 — FVP-7 stays visible. No entry may be reported as validated. The reproducibility gate is open on all nine and the report says so.

PF-9 — Run F-07's intervention test first if you are going to run only one test. It is the cheapest and most direct falsification in the formulary: remove `k` actions and check whether initiation multiplies by `exp(-bk)`.

PF-10 — Other people's ledgers are asked for, never assumed. Where an echo source is a person (F-08) or a cost is carried by a person (F-05, F-09's transfer audit), consent is obtained, a

decline path exists, and a decline is recorded as **unknown**, never as satisfied and never as a zero. *(CANON §5.4, binding on every entry.)*

PF-11 — Demote in print when a gate fails. A CONDITIONAL PASS whose condition is not met becomes a heuristic ordering in the report, in the same typeface as everything else, with the demotion stated. **A demoted formula is still useful; a formula that hides its demotion is not.**

PF-12 — Freeze before claiming. Any claim that the formulary works follows FVP-6 in full, including the shuffled-input control.

The Failure Modes

Formulary-wide faults — the ones that come from using the collection rather than from any single entry.

Fault	Diagnostic signature	Fix
Consistent-and-wrong	Every number agrees with every other; the picture is coherent and the practice does not behave as predicted	`M` is mismeasured. Six entries are fed by F-01, so a single upstream error propagates coherently. Re-run F-01 with an independent `E` estimate
Gate laundering	The FVP tables are present but every verdict is PASS	Sixty-three gates and no failures is not rigour; it is a copied table. Check FVP-2 and FVP-7 first, since those are the two that genuinely bind
Precision inflation	Four decimal places on a quantity estimated from three observations	PF-5. Report `n` and spread. The formulary's arithmetic is exact; its inputs are not
Formula shopping	Running entries until one produces a flattering number	Run in dependency order (PF-1) and report all nine, including the ones that did not help
Parameter drift	`N`, `delta`, `eta` quietly revised mid-block until the model fits	PF-2 and PF-4. Freeze, then measure. A parameter revised to fit is a parameter that has stopped measuring
Scope leakage	Internal-scope entries applied to a team without re-derivation	Volume Two's thesis is that the parameters change between scopes, not that the numbers transfer. Chapter 51 is the interpersonal formulary; use it
The single-instance conclusion	A whole regime redesigned from one bad week	Every entry's protocol specifies a minimum `n`. One instance is a prompt to measure, not a finding

Fault	Diagnostic signature	Fix
Cross-person comparison	"My `delta` is worse than hers"	FVP-7 is open on all nine. Within-person, within-domain only, until the reproducibility study exists
Ordinal contamination	Ratios and logs computed from ratings scales anywhere in the chain	FVP-2. Trace the chain: an ordinal input at F-01 voids F-02, F-06, F-08 and corrupts F-04, F-05, F-09
Moral arithmetic	The formulary used to argue that someone's `Q` matters less	CANON §5.4. The formulary measures; it does not rank persons, and any use that does has left the manual

---

Worked Example

One operator, one twelve-week block, all nine entries. The data are an **illustrative construction** used to demonstrate the full procedure; they describe no real person, and every derived number below was computed.

Setting. An operator running a first full formulary review on a domain of work they have been at for a year. Period = one week. Modulus = 50, assigned to a reference encounter from the preceding quarter and held for the block.

Step 1 — F-03, the ceiling

Framings elicited independently and in writing (Protocol step 1), then compared:

```
N     = 6                      ln N = 1.7918 nats = 2.5850 bits
u     = 0.38 (measured)
N_eff = 1 + 5(0.38) = 2.900    F_int = ln(2.900) = 1.0647 nats
claimed = 1.0647 / 1.7918 = 59.4% of the ceiling
```

Six framings audit to 2.9. That is the first finding of the review and it reframes everything after it: the operator experiences six perspectives and holds slightly under three.

Step 2 — F-01 and F-02, the moment and the calibration

wk	D	P	E	r	|ln r|	M	M_norm
1	62	41	55	1.5122	0.4136	83.171	1.2418
2	58	40	50	1.4500	0.3716	72.500	1.0825
3	66	44	58	1.5000	0.4055	87.000	1.2990
4	60	43	52	1.3953	0.3331	72.558	1.0834
5	64	46	60	1.3913	0.3302	83.478	1.2464
6	59	45	54	1.3111	0.2709	70.800	1.0571

wk	D	P	E	r	|ln r|	M	M_norm
7	55	47	57	1.1702	0.1572	66.702	0.9960
8	52	48	53	1.0833	0.0800	57.417	0.8573
9	50	49	59	1.0204	0.0202	60.204	0.8989
10	49	50	55	0.9800	0.0202	53.900	0.8048
11	51	52	58	0.9808	0.0194	56.885	0.8494
12	50	51	56	0.9804	0.0198	54.902	0.8198

M_norm uses a declared range [0, 120] scaled to [0, ln 6].

```
mean r,     weeks 1–4  = 1.4644      weeks 9–12 = 0.9904
mean |ln r|, weeks 1–4 = 0.3809      weeks 9–12 = 0.0199
mean M_norm over 12 weeks = 1.0197 nats   (ceiling 1.7918)
```

The operator arrived over-framed by about 46% and left calibrated. Distortion fell from 0.3809 nats — 0.55 bits, 1319 cents — to 0.0199 nats, which is below any plausible deadband. This is the review's success story, and Step 4 is where it turns.

Step 3 — F-08, the echo

Sent moment recorded before the return; returns from two independent sources, timestamped to the exchange:

wk	sent	back	rho	Delta_R
1	55	38	-0.3697	0.3697
2	58	42	-0.3228	0.3228
3	60	47	-0.2442	0.2442
4	57	46	-0.2144	0.2144
5	62	53	-0.1568	0.1568
6	60	54	-0.1054	0.1054
7	58	55	-0.0531	0.0531
8	61	60	-0.0165	0.0165
9	59	57	-0.0345	0.0345
10	60	61	+0.0165	0.0165
11	62	60	-0.0328	0.0328
12	61	62	+0.0163	0.0163

```
 mean rho,     weeks 1–4 = −0.2878      weeks 9–12 = −0.0086
mean Delta_R, weeks 1–4 =  0.2878    weeks 9–12 =  0.0250  sd(9–12) = 0.0289
 implied over-frame, weeks 1–4 = exp(0.2878) = 1.3335
 R_net = −1.5174   R_gross = 1.5830   E_R = 0.9586
```

Fitting the correction gain from the decay of |rho| over weeks 1–9:

```
per-week factor g = 0.6911     ->    lambda = 0.3089      R^2 = 0.8808
```

Three independent instruments agree. F-02 says the frame was over-set by 1.46 and came into calibration. F-08 says the returns were under by a factor implying an over-frame of 1.33 and

came into calibration. The fitted gain of 0.309 is in the stable monotone band, and `E_R = 0.96` confirms steady directed correction rather than thrash. **The four returns in weeks 9–12 fire the flattery alarm, and the scatter test clears it: the residuals change sign and their sd of 0.0289 is real variation, not a mirror.**

Step 4 — F-04, the storage law, and the turn

Cost from Chapter 5's ledger; `eta` fitted per F-04's protocol.

```
eta = 0.28      delta = 0.09      C = 0.30      Q_0 = 0.40
```

wk	M_norm	Q_t	step dQ	h_t	L_t	admissible?
0	1.2418	0.4000	+0.0117	0.1657	0.9418	yes
1	1.0825	0.4117	-0.0339	0.1646	0.7825	yes
2	1.2990	0.3778	+0.0297	0.1677	0.9990	yes
3	1.0834	0.4075	-0.0333	0.1650	0.7834	yes
4	1.2464	0.3742	+0.0153	0.1680	0.9464	yes
5	1.0571	0.3895	-0.0391	0.1666	0.7571	yes
6	0.9960	0.3504	-0.0527	0.1702	0.6960	yes
7	0.8573	0.2977	-0.0868	0.1749	0.5573	yes
8	0.8989	0.2110	-0.0673	0.1827	0.5989	yes
9	0.8048	0.1437	-0.0876	0.1888	0.5048	yes
10	0.8494	0.0561	-0.0672	0.1966	0.5494	**NO**
11	0.8198	-0.0111	-0.0695	0.2027	0.5198	**NO**
12	—	-0.0806	—	—	—	—

The operator got calibrated and went broke. Stored order fell from 0.40 to below zero over twelve weeks, and the admissibility condition `C <= (1−delta)Q + eta M` failed at week 11 — the week the model itself stops applying, which is a finding and not a rounding artifact.

Step 5 — F-09, the arrival that is a warning

```
mean M_norm = 1.0197
Q_inf = (0.28 * 1.0197 - 0.30) / 0.09 = -0.1609 Q-units       NEGATIVE
tau = 10.603 weeks      half-life = 7.350 weeks
```

Theorem 3.5's failure clause, live: `eta M* < c*`, so the fixed point is negative, and the loop converges anyway — toward the destruction of the thing. F-09's Protocol step 3 exists precisely to catch this **before** a practice is run; the operator ran it after, which is what a first review usually looks like and is why the review was worth doing.

The margin of failure is the number that makes this example worth reading:

```
maximum sustainable cost at u = 0.38:  C <= eta * F_int = 0.28 * 1.0647 = 0.2981
 observed cost                           :   C = 0.3000
 overshoot                              :   0.0019 Q-units per week  =  0.63%
```

The regime failed by 0.63 percent. Six-tenths of one percent over the sustainability ceiling of Corollary 3.4, sustained for twelve weeks, took the operator from 0.40 Q-units to below zero. There

was no crisis, no bad month, no identifiable mistake. **Corollary 3.4 is the single most important sociological statement in this manual, and this is its personal-scope face: past the ceiling, no choice of interaction geometry prevents order from decreasing, and the arithmetic does not care that the overshoot was small.**

Step 6 — F-05, the headroom paradox

Headroom was **positive and rising** through the entire collapse: 0.1657 at week 0, 0.2027 at week 11. At the fixed point, `h_inf = eta(ln N − F_int) = 0.28 ∗ (1.7918 − 1.0647) = 0.2036`.

Headroom rose because `Q` was falling, and `h` is reduced by `delta Q`. The gauge was correctly reporting that capacity for the new existed *at perfect geometry* while the actual geometry was losing 0.63% per week. This is F-05's central caveat in a live case: **`h > 0` is necessary and not sufficient, and an operator who watched only the headroom gauge would have seen an improving number for twelve weeks.** The geometry gap is the missing piece:

```
geometry gap = eta ( ln N − mean M_norm ) = 0.28 ∗ (1.7918 − 1.0197) = 0.2162
```

That 0.2162 is the part of the ceiling the operator's framings never reached, and it is why `h = 0.20` bought nothing.

Step 7 — F-06 and F-07, the two levers

```
F-06:  sum L over 12 weeks = 8.6365    mean L = 0.7197    L = 0 weeks: 0
       L_max = ln N − C = 1.7918 − 0.30 = 1.4918
```

The launch budget was healthy throughout — never zero, averaging 48% of its ceiling. F-06's "busy trap" in reverse: the operator had budget to imagine with and was losing stored order anyway. `L` and `dQ` are different quantities, and a good `L` does not imply a good `Q`.

F-07, actions-to-begin against weekly initiation count:

actions k	observed n	fitted n
1	22	23.19
2	15	14.86
3	10	9.52
5	4	3.91
8	1	1.03

```
ln n = 3.5888 − 0.4451 k      R^2 = 0.9989
halving distance = 1.557 actions
 removing 2 actions from the five-step practice: predicted count 3.91 −
> 9.52
```

Removing two actions from the front of the five-step practice is predicted to multiply its frequency by 2.43. This is the review's cheapest intervention and, per PF-9, the one whose prediction should be tested first.

Step 8 — the decision

Four candidate regimes, computed. `M_bar` under a raised `u` is scaled by the ratio of the new `F_int` to the old, which is a modelling choice and is stated as one:

Regime	u	C	F_int	M_bar	Q_inf	h_inf
as-is	0.38	0.30	1.0647	1.0197	**-0.1609**	0.2036
raise u	0.55	0.30	1.3218	1.2659	+0.6050	0.1316
cut C	0.38	0.24	1.0647	1.0197	+0.5058	0.2036
both	0.55	0.24	1.3218	1.2659	**+1.2716**	0.1316

The minimum requirements, computed exactly:

```
minimum sustained M for Q_inf >= 0:  M >= C/eta = 0.30/0.28 = 1.0714 nats  (59.8% of ln N)
  equivalent u, if M is read as F_int: u = (e^1.0714 − 1)/5 = 0.3839
  cost that would place Q_inf at 0.5:  C = eta∗M_bar − delta∗0.5 = 0.2405
```

The operator's measured u of 0.38 is 0.0039 below the u of 0.3839 that their cost requires. The whole twelve-week decline sits in that fourth decimal place, and both routes out of it are small: cut weekly cost by 0.06 Q-units, or raise u from 0.38 to 0.49.

Step 9 — the audit that closes the review

PF-10 and F-09's transfer audit. The operator's drive `eta F_int = 0.2981` per week is entirely their own — no transfer term. But part of their `C` is carried by a household member who absorbs displaced tasks. **That person's headroom is not the operator's to compute.** The review records the question as *asked, answer pending*, with the person free to decline, and the entry stands as **unknown** rather than as satisfied.

This is the only line in the review that cannot be closed by arithmetic, and CANON §5.4 is why it is in the report at all.

What the review found, in one paragraph

Six framings that audit to 2.9. A frame over-set by 46% that self-corrected to calibration in eight weeks at a stable gain of 0.31, confirmed independently by the echo log. A healthy launch budget throughout. A rising headroom gauge. And underneath all of it, a weekly cost 0.63% above the sustainability ceiling, converting twelve weeks of genuine improvement into a negative fixed point. No single entry would have found that: F-02 and F-08 said the operator was improving and were right; F-05 said capacity was available and was right. **Only F-04 and F-09, run together against a `Q_inf` computed rather than felt, showed that the room this life was converging on had less in it than the room it started from — by six-tenths of one percent.**

That is what a formulary is for.

Part IX — The Interpersonal Formulary

Chapter 46 . The Dyad

Part VIII took the instrument inward and measured one person. Part IX turns it outward — not back to Volume One's rooms and institutions, but to the thing that sits between two nervous systems, and to what quantity, exactly, is produced there.

We start at two because two is where sociology begins. One person is psychology. Three is already a network with three channels, a majority and a possible coalition. Two is the irreducible case: one relation, no majority, no topology. **The dyad is the atom of sociology, and everything in Volume One's Part IV is a compound of it.**

It is also the case in which the surplus law is stated in its original form. Paper A, Theorem 2.1 is a statement about a *pair* of states. Every application to a group of forty is a modelling extension of a two-body result, badged in Chapter 7 as the Roster Reduction. In the dyad there is no extension and no badge to apply.

That is why this chapter contains the sharpest single statement in the volume, and why it is the oldest. Theorem 4.3(a) says the surplus of a pair is zero if and only if their states are parallel. In the older canon of this house the same sentence reads: **the wedding of like with like is a non-event.** One is a theorem about unit vectors; the other is tradecraft that predates the theorem by a long way. This chapter shows they are the same claim, computes the price of ignoring it, and specifies what to build instead.

The Law

46.1 The pair surplus, stated exactly

From Paper A, Theorem 2.1, for two unit states x and y in a space that can express N distinct categories:

```
F(x,y) = f( u(x,y) )

u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2        in [0, 1]

f(u)   = ln( 1 + (N−1) u )                uniquely characterised by S1–S4

N_eff  = e^{F} = 1 + (N−1) u               the effective count of the pair
```

For unit states `x . y = cos θ`, and therefore the identity that runs this chapter:

```
u = 1 − cos^2 θ = sin^2 θ
```

u is the squared sine of the angle between two people's states. It is zero when they point the same way, zero when they point exactly opposite ways, and one when they are at right angles. That last clause is not a rhetorical flourish and it is the reason §46.10 exists.

46.2 Theorem 4.3(a), and the old canon

> **Paper A, Theorem 4.3(a) [PROVEN].** `F(x,y) = 0` if and only if x and y are parallel (`(x . y)^2 = 1`, equivalently `u = 0`).

The "only if" direction does the work. It is not merely that parallel states produce little; it is that **zero surplus occurs in exactly one circumstance and no other.** If a pair is producing nothing, the law names the cause uniquely — no second explanation exists inside the formalism. Either the states are parallel, or the surplus is strictly positive.

The house canon states the same fact in the older register:

> **THE WEDDING OF LIKE WITH LIKE IS A NON-EVENT.**

Two things must be said about that sentence. First, the **mathematical content is proven**: `f(0) = ln(1 + (N−1).0) = ln 1 = 0`, and by strict monotonicity (Thm 3.1) nothing else reaches zero. Second, the **sociological reading is a modelling proposal, badged**:

> **THE DYADIC IDENTIFICATION [POSTULATE].** We take the state x of a person, with respect to a given decision domain, to be a unit vector in an N-category space whose components are that person's weighting across the domain's N distinguishable categories. Two people's non-parallelism u is then estimable from disagreement, error-decorrelation, and prediction diversity by the protocols of Chapter 9 and Appendix C.

Paper A does not say human beings are unit vectors. It says that *if* a system's interaction surplus depends on pair geometry through S1-S4, *then* it takes exactly one form. The identification of a person with a state is ours, is testable (Ch 25), and is badged here so that no reader mistakes the theorem's certainty for the identification's.

46.3 What N is in a dyad, and what it is not

This is the most common error made with the pair formula, and it must be closed first.

N is not the headcount. The headcount in a dyad is two, always; if N were the headcount the ceiling on pair surplus would be `ln 2 = 0.6931` nats for every pair in the world. N is the **number of distinguishable categories the ambient domain can express** — the dimension of the space the two states live in — entering through normalisation axiom S4, `f(1) = ln N`.

A couple deciding how to spend a Saturday works in perhaps `N = 6` (rest, repair, obligation, pleasure, money, others). A two-person engineering partnership deciding a structural approach may work in `N = 12` or more. **Two people can carry a large N between them; that is what makes a good pair valuable.**

```
ln N is the hard ceiling on the surplus of ONE interaction   (Paper D, Cor 2.2)

 N =  2  ->  ln N = 0.6931
 N =  4  ->  ln N = 1.3863
 N =  6  ->  ln N = 1.7918
 N =  8  ->  ln N = 2.0794
 N = 12  ->  ln N = 2.4849
 N = 16  ->  ln N = 2.7726
 N = 24  ->  ln N = 3.1781
```

A pair working an `N = 12` domain at right angles produces 2.4849 nats; the same two people, parallel, produce 0.0000 in the same domain. **The domain sets the ceiling; the angle decides how much of it you get.**

46.4 The honesty clauses, restated at dyad scope

1. **Paper A is pure mathematics and makes no empirical or metaphysical claim.** Theorem 4.3(a) is about vectors. That partnerships behave this way is this manual's proposal, and Chapter 55 says what would kill it.
2. **As deployed, the surplus functional is a lower bound** on the pair's information, not the information itself. An `N_eff` of 3.09 does not mean the couple contains exactly 3.09 perspectives. It means the geometry we managed to measure guarantees at least that much. We log the floor.
3. **Surplus is not Shannon entropy or mutual information** (Paper D). The `ln` is not a Shannon `ln`. Do not import channel-capacity intuitions.
4. **The triangle inequality FAILS for geometric triples** (Paper A §5). This is the one honesty clause a dyad chapter uses as a positive result: because surplus cannot be chained across three states, **the pair is the only scale at which the law is exact** — an argument for measuring pairs directly rather than inferring them from a group figure.
5. **S3 is a design axiom.** `N_eff = 1 + (N−1)u` is affine in u because we chose it to be. Every linear-looking result in this chapter is downstream of that choice.

46.5 The dyadic storage law

The pair does not merely produce surplus. It stores it. Paper E §3, unchanged in form:

```
Q_{t+1} = (1 − δ) Q_t + η S_t − C_t          S_t = ln(1 + (N−1) U_t)
```

At dyad scope the four parameters read:

Symbol	In a pair
`Q`	the shared store: trust, shorthand, joint method, common references, the built thing
`δ`	drift, forgetting, absence, the erosion of shorthand when unused
`η`	how much of what passes between them becomes durable rather than merely pleasant
`C`	maintenance: coordination, repair, the cost of being known

Steady state (Thm 3.5): `Q_inf = (η F(u∗) − c∗) / δ`. **Headroom (Paper E §6.4):** `h_t = η ln N − δ Q_t − C_t`.

Both apply without modification. The dyad is not a special case of the storage law; it is the storage law at `n = 2`.

The Engineering

46.6 Angle is the control variable

Because `u = sin^2 θ`, every statement about a pair can be made in degrees. This is the most useful translation in the chapter, because people can estimate an angle and cannot estimate a

squared wedge norm.

θ (deg)	cos θ	u = sin^2 θ
0	+1.0000	0.0000
5	+0.9962	0.0076
10	+0.9848	0.0302
15	+0.9659	0.0670
20	+0.9397	0.1170
25	+0.9063	0.1786
30	+0.8660	0.2500
40	+0.7660	0.4132
45	+0.7071	0.5000
50	+0.6428	0.5868
60	+0.5000	0.7500
70	+0.3420	0.8830
80	+0.1736	0.9698
90	+0.0000	1.0000

Read the top of that table carefully. **A pair fifteen degrees apart has `u = 0.067` — seven per cent of the way to orthogonal, not seventeen.** The squared sine punishes small angles severely: two colleagues who "see things a bit differently" are, in the law's accounting, almost the same instrument. The other end is equally instructive — between 60 and 90 degrees u moves only from 0.75 to 1.00. **Most of the available non-parallelism is bought in the first sixty degrees.** You do not need a partner who is your opposite in everything, only one who is genuinely oblique.

46.7 The pair-surplus table

`F(u) = ln(1 + (N−1)u)`, computed across the domain sizes a real pair might work in.

u	N=2	N=3	N=4	N=6	N=8	N=12	N=16	N=24
0.00	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
0.01	0.0100	0.0198	0.0296	0.0488	0.0677	0.1044	0.1398	0.2070
0.05	0.0488	0.0953	0.1398	0.2231	0.3001	0.4383	0.5596	0.7655
0.10	0.0953	0.1823	0.2624	0.4055	0.5306	0.7419	0.9163	1.1939
0.20	0.1823	0.3365	0.4700	0.6931	0.8755	1.1632	1.3863	1.7228
0.30	0.2624	0.4700	0.6419	0.9163	1.1314	1.4586	1.7047	2.0669
0.50	0.4055	0.6931	0.9163	1.2528	1.5041	1.8718	2.1401	2.5257
0.70	0.5306	0.8755	1.1314	1.5041	1.7750	2.1633	2.4423	2.8391
0.90	0.6419	1.0296	1.3083	1.7047	1.9879	2.3888	2.6741	3.0773
1.00	0.6931	1.0986	1.3863	1.7918	2.0794	2.4849	2.7726	3.1781

And the effective count, which is the same information in the units people can actually argue about:

u	N=4	N=8	N=12	N=24
0.00	1.000	1.000	1.000	1.000
0.05	1.150	1.350	1.550	2.150

0.10	1.300	1.700	2.100	3.300
0.20	1.600	2.400	3.200	5.600
0.30	1.900	3.100	4.300	7.900
0.50	2.500	4.500	6.500	12.500
0.70	3.100	5.900	8.700	17.100
0.90	3.700	7.300	10.900	21.700
1.00	4.000	8.000	12.000	24.000

Sit with one fact in that second table. In an N = 12 domain a pair at u = 0.50 has an effective count of **6.5** — two people, six and a half. The same pair at u = 0.05 has an effective count of **1.55**.

The same two human beings are worth 6.5 or 1.55 depending only on the angle between their states. Talent does not move that number; talent lives in N and η. The angle is its own parameter, and it is the one usually left unmanaged.

46.8 The marginal structure — where the value actually sits

Differentiating: dF/du = (N−1) / (1 + (N−1)u). At N = 12:

u	F(u)	dF/du
0.00	0.0000	11.0000
0.05	0.4383	7.0968
0.10	0.7419	5.2381
0.20	1.1632	3.4375
0.30	1.4586	2.5581
0.50	1.8718	1.6923
0.70	2.1633	1.2644
0.90	2.3888	1.0092
1.00	2.4849	0.9167

The marginal return on non-parallelism at u = 0 is **twelve times** the marginal return at u = 1. Strict concavity (Thm 3.3) again, and per CANON_INTERNAL §2 this is the lens, not pessimism: a logarithm looks like diminishing returns from inside.

The engineering consequence is asymmetric and worth stating hard. **The first small increment of difference in a pair is the most valuable thing you will ever add to it, and the first small increment of convergence is the most expensive thing you will ever lose.** A pair that drifts from u = 0.10 to u = 0.05 loses 0.3036 nats over an interval of 0.05 — a rate of 6.07 nats per unit of u. A pair that drifts from u = 0.90 to u = 0.50 loses 0.5170 nats over an interval of 0.40 — a rate of 1.29 nats per unit. **Difference lost near the parallel wall is worth nearly five times what the same difference is worth near orthogonality.** Near-parallel pairs sit on the steep part of the curve, which is where small maintenance efforts pay best and small neglects cost most.

As a fraction of the ceiling at N = 12, F/ln N runs 0.1764 at u = 0.05, 0.2986 at 0.10, 0.4681 at 0.20, 0.5870 at 0.30, 0.7533 at 0.50, 0.8706 at 0.70 and 0.9613 at 0.90. **A pair at 45 degrees (u = 0.50) already captures three-quarters of what a perfectly orthogonal pair could.** Perfection is not the target. Obliqueness is.

46.9 Opposition is also parallel — the result nobody expects

Write the alignment between the two states as ρ = x . y, a correlation in [-1, +1]. Then u = 1 − ρ^2, and u depends only on the *magnitude* of ρ. At N = 12:

ρ	u	N_eff	F
+1.00	0.0000	1.0000	0.0000
+0.95	0.0975	2.0725	0.7288
+0.90	0.1900	3.0900	1.1282
+0.70	0.5100	6.6100	1.8886
+0.50	0.7500	9.2500	2.2246
0.00	1.0000	12.0000	2.4849
−0.50	0.7500	9.2500	2.2246
−0.90	0.1900	3.0900	1.1282
−1.00	0.0000	1.0000	0.0000

Look at the last row. **A pair in perfect opposition has u = 0, N_eff = 1, and surplus exactly zero — the same as a pair in perfect agreement.** Theorem 4.3(a) covers both cases, because it is stated on (x . y)^2 = 1, which is satisfied by ρ = +1 and by ρ = −1 alike.

This is not a quirk of notation but a correct sociological statement following from the geometry with no interpretive help:

> **The mirror and the mirror-image are the same instrument.** A partner who agrees with you about everything and one who disagrees about everything carry identical information, because either position is fully recoverable from yours. Reflexive contradiction is a copy with a sign flip.

Much of what is called "creative tension" in partnerships is ρ ∼ −1 and generates no surplus at all. It generates heat, which enters the model as C, not S. **An opposition pair is a parallel pair with a cost overhead** — it pays C and returns zero surplus, which is a worse position than working alone, since a lone operator at least pays nothing for the redundancy.

This result is not local to the dyad, and it belongs in two other places in the manual. It is the same arithmetic as Chapter 7's roster reduction: u enters N_eff = 1 + (N−1)u and u sees only |ρ|, so a systematic contrarian is audited out of the effective count exactly as a systematic agreer is. And it is a **failure mode distinct from the Parallel Wall of Chapter 11**, sharing that fault's arithmetic while presenting as its opposite: the Parallel Wall is a room that sounds like consensus, the opposition pair is a room that sounds like debate, and both read N_eff = 1. **The contrarian who inverts you carries the same information as the sycophant who echoes you: none.** A diagnostic built to catch the Parallel Wall by listening for agreement will not catch this one, which is why Chapter 11 gives it its own catalogue entry and why the obliqueness test of §46.10 below is stated on predictability rather than on assent.

46.10 Complementarity is orthogonality, not opposition

The distinction the previous section forces is the practical centre of this chapter. Reading the ρ table by name: ρ = +1.00 is identity, a copy; ρ = +0.95 is the culture-fit hire, N_eff 2.07; ρ = +0.50 is **complementarity**, N_eff 9.25; ρ = 0 is maximal complementarity, N_eff 12.00; ρ = −0.95 is the sparring partner, N_eff 2.07 again; ρ = −1.00 is identity inverted, also a copy.

Complementarity is a middle condition, and both extremes of agreement are its failure modes. Advice that says "find someone who completes you" is, read through the law, correct; advice that says "find someone who challenges everything you say" prescribes the right-hand collapse of the same curve.

The operational test is not "do we agree?" and not "do we argue?". It is:

> **THE OBLIQUENESS TEST.** Given the partner's position on a question in the domain, can you predict it from your own — by copying it, or by inverting it? If either rule works reliably, u is near zero and the pair is producing nothing. **A useful partner is one whose position you cannot compute from yours by any simple rule.**

46.11 Crossing and divergence inside a pair

Theorem 4.1 decomposes non-parallelism into two sources, and it applies to a pair as cleanly as to a roster. For a pure block vector:

```
u = u_cross + u_div      u_cross = β^2        (between-category difference)
                         u_div   = α^2 sin^2 γ  (within-category difference)
                              α^2 + β^2 = 1

 max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

At N = 12, worked:

β^2	γ	u_cross	u_div	u	f(u_c)	f(u_d)	f(u)	max	sum
0.00	30deg	0.0000	0.2500	0.2500	0.0000	1.3218	1.3218	1.3218	1.3218
0.25	30deg	0.2500	0.1875	0.4375	1.3218	1.1192	1.7600	1.3218	2.4410
0.25	60deg	0.2500	0.5625	0.8125	1.3218	1.9723	2.2963	1.9723	3.2941
0.50	45deg	0.5000	0.2500	0.7500	1.8718	1.3218	2.2246	1.8718	3.1936
0.60	20deg	0.6000	0.0468	0.6468	2.0281	0.4152	2.0937	2.0281	2.4434
0.00	80deg	0.0000	0.9698	0.9698	0.0000	2.4569	2.4569	2.4569	2.4569

Every row satisfies `max <= f(u) <= sum`, as Theorem 4.1 requires, and the bound is never tight in the middle rows — the honest content of subadditivity: **you never get more than the sum of the two sources, and usually noticeably less.** No synergy beyond additivity is promised, and a design that assumes it is unfunded.

At dyad scope the two sources have plain names:

- **`u_cross` — the pair spans two trades.** An engineer and a clinician. A composer and a statistician. Different vocabularies, different training lineages, different failure intuitions.
- **`u_div` — the pair shares a trade and disagrees inside it.** Two structural engineers who differ on the same question. Same vocabulary, different position.

They are not substitutes. The β^2 = 0.60, γ = 20 deg row is the cross-disciplinary pair that never argues; the β^2 = 0, γ = 30 deg row is the same-trade pair that argues constantly. Both have real surplus. Neither has both kinds. **Diagnose both, and do not accept one as evidence of the other.**

46.12 The dyad's cost floor, and why pairs are the cheapest structure in existence

Brooks' count of communication channels, n(n−1)/2:

n	channels	channels per person
2	1	0.50
3	3	1.00
4	6	1.50
5	10	2.00
6	15	2.50
8	28	3.50

The dyad is the only human structure with fewer channels than people. Everything else costs more in coordination than it has heads to pay with. That is the structural reason pair-based practices — the two-person cockpit, the two-person research collaboration, the marriage — recur across every domain that works under load.

It also means the dyad has the highest headroom per person available to any social form, because C is at its floor: Volume One's Chapter 5 crossover calculation, read from this side, is the observation that the dyad is on the good side of every crossover.

The corollary is not comfortable. **A pair that produces nothing has no excuse available.** It cannot blame coordination overhead, because it has the minimum. Under Theorem 4.3(a) there is exactly one remaining cause.

46.13 Reciprocity: u is symmetric, η is not

u(x,y) = u(y,x) identically — it is built from (x . y)^2, which is symmetric. So:

> **The surplus of a pair is a joint property. Neither person owns it, neither person produced it, and it cannot be divided between them.**

No accounting scheme can attribute a pair's surplus to one member, and any incentive structure that tries is measuring something the law says does not exist.

Conversion is a different matter, and conversion is where reciprocity lives. Let η_{A->B} be the rate at which what A offers becomes durable structure in B, and η_{B->A} the reverse. Nothing requires these to be equal.

> **THE DIRECTIONAL CONVERSION POSTULATE [POSTULATE].** A dyad has two conversion parameters, one per direction. The joint store converts at the mean (η_{A->B} + η_{B->A}) / 2; each partner's private store converts at the rate directed toward them. **Reciprocity is the condition η_{A->B} = η_{B->A}, and its violation is measured on a log-ratio:**

reciprocity asymmetry a = | ln(η_{A->B} / η_{B->A}) |

a = 0	symmetric
a = 0.4055	a 3:2 imbalance
a = 0.6931	a 2:1 imbalance
a = 1.3863	a 4:1 imbalance

```
a = 2.1972   a 9:1 imbalance
```

Now run each partner's private store at δ = 0.15, C_A = C_B = 0.25, u = 0.45, N = 12, so S = 1.7834, holding the **mean conversion fixed at 0.50** in every row:

η_A->B	η_B->A	mean	a	Q_A(inf)	Q_B(inf)	sum
0.50	0.50	0.50	0.0000	+4.2780	+4.2780	+8.5559
0.60	0.40	0.50	0.4055	+3.0890	+5.4669	+8.5559
0.70	0.30	0.50	0.8473	+1.9001	+6.6558	+8.5559
0.80	0.20	0.50	1.3863	+0.7112	+7.8448	+8.5559
0.90	0.10	0.50	2.1972	−0.4777	+9.0337	+8.5559

The sum is constant down the whole column and the individual outcomes are not — the thing any aggregate measurement will miss:

> **Reciprocity failure is invisible in the joint account and total in the private one.** The partnership's books look identical at a = 0 and a = 2.20. At a = 2.20 one partner's private store has a *negative* steady state — Paper E Theorem 3.5's declared failure mode, reached not by any deficiency of the pair's geometry but purely by the direction of its conversion.

A partnership called "productive but one-sided" is describing exactly this. The productivity is real and sits in the sum; the one-sidedness is real and sits in the split. Both are true at once.

46.14 The return difference

Volume Two's fourth derived formula (CANON_INTERNAL §8), stated for a pair. M is what one partner sends into the relation; M' is what comes back.

```
Δ_R = | ln( M' / M ) |
```

M'/M	Δ_R	
3.00	1.0986	
2.00	0.6931	
1.50	0.4055	
1.25	0.2231	
1.10	0.0953	
1.00	0.0000	<- no information returned
0.90	0.1054	
0.75	0.2877	
0.50	0.6931	
0.33	1.1087	

Δ_R = 0 means the pair returned your own signal unchanged. **An echo is not a reply.** In the dyad this is the behavioural face of u = 0: the partner whose response is always your own position restated is not a partner, they are a surface. The magnitude tells you how much came back; the sign of ln(M'/M) tells you which direction the world moved you.

46.15 Every working pair destroys its own u

Here is the structural problem that makes this a maintenance chapter, not a selection chapter.

A pair that works together converges: shared vocabulary, shared references, shared priors, shared blind spots. Every one of those is Q going up and u going down. **The dyad's stored order and its input geometry are in direct conflict, and the conflict is not a design error. It is what success looks like from the inside.**

Model the drift as geometric decay, $u_t = u_0 (1 - \lambda)^t$, and the half-life follows:

λ	half-life of u (periods)
0.02	34.31
0.05	13.51
0.08	8.31
0.10	6.58
0.20	3.11
0.30	1.94

At $u_0 = 0.60$, $\lambda = 0.10$, $N = 12$, twenty periods run as follows:

t	u	N_eff	F
0	0.6000	7.6000	2.0281
5	0.3543	4.8972	1.5887
10	0.2092	3.3013	1.1943
15	0.1235	2.3589	0.8582
20	0.0729	1.8024	0.5891

Twenty periods of successful collaboration take a pair from an effective count of 7.6 to an effective count of 1.8. Nobody did anything wrong. They agreed with each other, repeatedly, because they were learning from each other, which is what they were supposed to do.

This is why Chapter 11 says u is not a property you acquire but a quantity that decays. In the dyad the point is sharpest, because there is no third party to re-inject difference. **A pair has no internal source of u. It must import.**

46.16 The minimum viable angle

Set the steady state to zero and solve for the u below which the pair holds no store at all. $Q_{inf} >= 0$ requires $\eta\ F(u) >= C$, so:

$$S_{min} = C / \eta$$
$$u_{min} = (e^{C/\eta} - 1) / (N - 1)$$

For a pair with $\eta = 0.55$, $C = 0.34$, $N = 12$:

$$S_{min} = 0.34 / 0.55 = 0.6182$$
$$u_{min} = (e^{0.6182} - 1) / 11 = 0.0778$$
$$N_{eff} \text{ at } u_{min} = 1.8556$$

Below $u = 0.078$ this pair sustains no joint store whatsoever. The relationship does not merely become less productive; its equilibrium goes negative and everything it built erodes at

rate δ. In degrees, `u = 0.0778` is `θ = 16.2`. **Sixteen degrees is this partnership's survival line**, and every pair has such a line, computable from three numbers it can measure.

46.17 Maintenance beats selection — computed

If u decays and there is no internal source, the only lever left is periodic re-injection: importing difference on a schedule. Volume One's Chapter 2 result applies — **a leaky integrator responds to cadence, not intensity** — and here it can be costed.

Take the same pair (`η = 0.55`, `δ = 0.12`, `C = 0.34`, `N = 12`) starting at `u_0 = 0.62` with drift `λ = 0.06` per quarter, and compare restoring u to 0.62 every k quarters against no restoration, measured by the store after 24 quarters:

```
restoration interval    Q at t = 24
--------------------    -----------
never                      2.7064
every 12 quarters          4.7070
every  8 quarters          5.3283
every  6 quarters          5.6209
every  4 quarters          5.8991
every  2 quarters          6.1615
```

Restoring once every twelve quarters — one deliberate importation of difference every three years — recovers **58 per cent** of the gap between doing nothing and doing it every six months (4.7070 vs 2.7064 against a maximum of 6.1615). Halving the interval to six quarters buys 0.914 further units; halving again buys 0.541.

The returns on restoration frequency are themselves concave, and the first restoration is worth more than all subsequent increases in frequency combined. A pair that never restores loses two-thirds of its potential store and experiences the loss as fatigue. What counts as a restoration is specified in S-12: not a holiday, but the deliberate introduction of a state the pair cannot predict from its own.

46.18 The formulas of this chapter, with their FVP tables

Every formula in Volume Two carries the full seven-gate Formula Validation Protocol (CANON_INTERNAL §6). This chapter's follow.

FVP TABLE 46-A — Non-parallelism `u = 1 − (x.y)^2`

Gate	Finding
FVP-1 DIMENSIONS	x, y unit vectors; `x.y` dimensionless; u dimensionless in [0,1]. Both sides agree. **PASS.**

Gate	Finding
FVP-2 SCALE TYPE	u is ratio-scaled with a true zero (exact parallelism) and a true maximum. The *components* of x need at least interval scale to normalise; with ordinal rankings only, u must come from a rank-correlation surrogate and **demotes to an ordering, not a magnitude. CONDITIONAL PASS — the encoder must be declared.**
FVP-3 LIMITING CASES	`u -> 0`: identical or exactly opposite states (§46.9). `u -> 1`: orthogonal, no mutual predictability. Both nameable; u is bounded, so no divergence. **PASS.**
FVP-4 FIXED POINTS	`u = 0`, the invariant of the copy relation and the unique zero of `F` (Thm 4.3a); `u = 1`, the invariant of independence. Both attained. **PASS.**
FVP-5 IDENTIFIABILITY	u recoverable from disagreement rates and error decorrelation (Ch 9, App C) without observing x, y. It does **not** distinguish `ρ = +0.9` from `ρ = −0.9`; sign information is lost and must be recovered separately. **PASS WITH DECLARED AMBIGUITY.**
FVP-6 FALSIFICATION	Preregister u from pre-discussion disagreement; predict held-out joint output; beat persistence, direct-input, matched autoregressive, shuffled-input. **Kill condition: u at `t` fails to improve prediction of `t+1` output over the shuffled control.**
FVP-7 REPRODUCIBILITY	The cross-pair test: same operationalisation, different observer, different pair, same domain — does u recover? **UNTESTED. This is the gate that decides whether the dyadic identification is science, and it has not been run.**

FVP TABLE 46-B — Pair surplus `F = ln(1 + (N−1)u)` and `N_eff = 1 + (N−1)u`

Gate	Finding
FVP-1 DIMENSIONS	`(N−1)u` dimensionless, so the logarithm's argument is dimensionless as required. `F` in nats; `N_eff` a pure count. **PASS.**

Gate	Finding
FVP-2 SCALE TYPE	N a count (ratio scale, true zero); u ratio-scaled per 46-A; F in nats inherits ratio scale. **PASS, conditional on 46-A's encoder.**
FVP-3 LIMITING CASES	`u -> 0`: `F -> 0`, `N_eff -> 1`, the non-event. `u -> 1`: `F -> ln N`, the ceiling. `N -> inf` at fixed `u > 0`: `F -> inf` logarithmically — real divergence, but so slow as to be useless, which is Chapter 8's warning. `N -> 1`: `F -> 0` for all u, correct, since a one-category domain admits no difference. **PASS.**
FVP-4 FIXED POINTS	`F = 0 <=> u = 0` (Thm 4.3a); `F = ln N <=> u = 1` (Thm 4.3b). Unique both ways. **PASS — the strongest fixed-point result in the volume.**
FVP-5 IDENTIFIABILITY	N from enumerating declared categories, u from the decorrelation protocol — separately recoverable. But F alone does **not** identify `(N, u): F = 0.6931` arises from `(N=2, u=1)` and from `(N=12, u=0.0909)`. **Report N and u, never F alone. PASS ONLY IF BOTH ARE REPORTED.**
FVP-6 FALSIFICATION	As 46-A, plus the preregistered prediction that output scales with `ln N_eff` rather than headcount or mean ability. **Kill condition: headcount or mean ability predicts held-out joint output better than `ln N_eff`.**
FVP-7 REPRODUCIBILITY	**UNTESTED at dyad scope.** Woolley 2010 is the nearest published evidence that group output tracks something other than member ability, but it is a group-scale finding in short-task laboratory settings, not a dyadic field result. Supporting context, not confirmation.

FVP TABLE 46-C — Dyadic storage `Q_{t+1} = (1−δ)Q_t + ηS_t − C_t`

Gate	Finding
FVP-1 DIMENSIONS	Q in declared order-units; S in nats; **η therefore carries units of [order]/[nat]** and must be reported with them; δ dimensionless per period; C in order-units per period. Both sides in order-units. **PASS, provided η's units are declared.**

Gate	Finding
FVP-2 SCALE TYPE	Requires Q ratio-scaled with a true zero. For a pair the order proxy — shared method, joint artefacts, documented decisions — must be countable. **PASS where the proxy is a count; FAILS and demotes where it is a satisfaction rating.**
FVP-3 LIMITING CASES	δ –> 0: no fixed point, Q_t <= Q_0 + tη ln N (Thm 3.3) — the sealed pair that never forgets and never revises. δ –> 1: total amnesia each period. η –> 0: the pleasant partnership that builds nothing. C large: negative Q_inf, Thm 3.5's declared failure. **PASS — every divergence names a real condition.**
FVP-4 FIXED POINTS	Q_inf = (ηF(u∗) − c∗)/δ for δ > 0, unique, reached geometrically by Banach since |1−δ| < 1. **PASS.**
FVP-5 IDENTIFIABILITY	δ from decay under zero input; η from attributable ΔQ per unit S; C from tally. **Known hazard (Paper E §6.3): if the order proxy Y repackages the input U, η and the surplus term are not separately identifiable.** Acute for pairs — "we have a great shorthand" is both. **CONDITIONAL PASS; the proxy must be independently justified.**
FVP-6 FALSIFICATION	Fit Y_{t+1} = aY_t + bF(U_t) − Ĉ_t + ε_t, freeze, report held-out, four baselines. **Kill condition: the shuffled-input control matches the fitted model held-out.**
FVP-7 REPRODUCIBILITY	**UNTESTED.** Do (δ, η) from one pair transfer to another pair in the same trade? Nobody has run it.

FVP TABLE 46-D — Reciprocity asymmetry a = \|ln(η_{A–>B} / η_{B–>A})\|

Gate	Finding
FVP-1 DIMENSIONS	Ratio of two identically-united quantities; argument dimensionless; a in nats. **PASS.**
FVP-2 SCALE TYPE	Requires both conversions ratio-scaled with a true zero (a direction that converts nothing). Attainable where η is attributable durable output per unit input. **PASS.**

Gate	Finding
FVP-3 LIMITING CASES	`a -> 0`: perfect reciprocity. `a -> inf`: one direction converts nothing — the audience relation, the extractive partnership. Both nameable. **PASS.**
FVP-4 FIXED POINTS	`a = 0`, which is exactly the reciprocity condition. It is **not** an attractor: nothing in the model drives a down, which is why reciprocity needs a protocol (S-14). **PASS, fixed point unstable in practice.**
FVP-5 IDENTIFIABILITY	The two directional η are identifiable only if each private store is measured separately. **From the joint store alone they are NOT identifiable** — every row of §46.13's table sums the same. **PASS ONLY UNDER SEPARATE MEASUREMENT; this is why "the partnership is doing well" is not evidence of reciprocity.**
FVP-6 FALSIFICATION	Preregister a from directional attribution; predict divergence of the two private stores. **Kill condition: a fails to predict store divergence beyond chance.**
FVP-7 REPRODUCIBILITY	**UNTESTED.**

FVP TABLE 46-E — Return difference `Δ_R = |ln(M'/M)|`

Gate	Finding
FVP-1 DIMENSIONS	`M` and `M'` share units by construction, so the ratio is dimensionless and `Δ_R` is in nats. **PASS.**
FVP-2 SCALE TYPE	**The binding gate.** `M'/M` requires `M` ratio-scaled with a true zero (Ch 53). With only ordinal reports, `Δ_R` demotes to "the reply differed / did not differ" — an ordering, not a magnitude. **CONDITIONAL, and in most dyadic settings the condition is not met without a magnitude-estimation procedure.**
FVP-3 LIMITING CASES	`M' -> M`: `Δ_R -> 0`, the echo. `M' -> 0`: `Δ_R -> inf`, total absorption. `M -> 0`: undefined, correctly, since nothing was sent. **PASS.**
FVP-4 FIXED POINTS	`M' = M`: the relation returns what entered it. **This is a null, not an ideal.** The design target is `Δ_R` bounded away from zero, not minimised. **PASS.**

Gate	Finding
FVP-5 IDENTIFIABILITY	Only the ratio is recoverable from Δ_R, never M and M' separately. Report ratio with direction. **PASS WITH DECLARED LIMITATION.**
FVP-6 FALSIFICATION	Preregister that pairs with persistently low Δ_R show lower held-out joint output. **Kill condition: no relation beyond the shuffled control.**
FVP-7 REPRODUCIBILITY	**UNTESTED, and blocked upstream by FVP-2.** Until magnitude estimation of M is established (Ch 38, 53) this gate cannot be attempted. Stated rather than skipped.

FVP TABLE 46-F — Non-parallelism decay `u_t = u_0 (1−λ)^t`

Gate	Finding
FVP-1 DIMENSIONS	λ dimensionless per period; u dimensionless; t a count. **PASS.**
FVP-2 SCALE TYPE	Ratio scale in u (46-A), count scale in t. **PASS.**
FVP-3 LIMITING CASES	λ −> 0: u constant — a pair that works together and never influences each other, which is the mark of a pair not actually collaborating. λ −> 1: collapse to u = 0 in one period, total capture of one partner by the other. Both real. **PASS.**
FVP-4 FIXED POINTS	u = 0, unique and **globally attracting for any λ > 0.** The chapter's central engineering fact: convergence is the default, difference is the maintained condition. **PASS.**
FVP-5 IDENTIFIABILITY	λ = 1 − (u_1/u_0)^{1/Δt} from two clean u measurements with no restoration between. **PASS.**
FVP-6 FALSIFICATION	Preregister λ from periods 1-8, predict u at period 16, compare to a constant-u baseline. **Kill condition: constant-u predicts held-out u as well as the decay model.**
FVP-7 REPRODUCIBILITY	**UNTESTED.** Geometric decay is the simplest assumption, not a measured law; pairs may converge in bursts around shared projects. **The functional form is a modelling choice, badged as one.**

46.19 The refusal clause

CANON §5.4 governs this chapter with unusual force, because dyadic interventions act on people who did not attend the design meeting. **Nothing here licenses arranging a person's partnerships for them, selecting their associates by measured obliqueness, or scoring a relationship they did not agree to have measured.** These instruments are for a pair to point at itself, or for an institution to point at partnerships it is itself constituting — a review pairing, a co-authorship, a two-person watch. The moment a protocol removes a person's ability to decline the measurement or the pairing, it is out of scope for this manual and stays out.

The Specification

Requirements for constituting and maintaining a two-person working relation. Each names the parameter it moves.

S-1 . Declare the domain and count it. Write down the domain's `N` — the distinguishable categories the pair's work must span — and justify each, before any pairing is assessed. `[N]` *Test: an independent reader lists the same categories from the written definition.*

S-2 . Never report `F` without `N` and `u`. Per FVP 46-B/FVP-5, `F` alone does not identify the pair. `[reporting]`

S-3 . Estimate `u` before the pairing is fixed. Independent pre-discussion estimates on at least eight domain questions; disagreement and error decorrelation per Chapter 9. `[u]` *Test: estimates recorded before the two people spoke about the questions.*

S-4 . Apply the Obliqueness Test. If either position is reliably recoverable from the other by copying or by inversion, record $u \sim 0$ and do not assume surplus. `[u]` *Test: predictive accuracy of the copy rule and of the inversion rule, both measured.*

S-5 . Reject opposition as a substitute for complementarity. A pair whose disagreement is systematic and total is $\rho \sim -1$, $u \sim 0$. Record it as a similarity failure, not as productive tension. `[u, C]`

S-6 . Measure `u_cross` and `u_div` separately and report both; a pair rich in one and starved of the other is a diagnosable condition with a specific fix. `[u]`

S-7 . Compute the minimum viable angle. `u_min = (e^{C/η} − 1)/(N−1)`; every standing pair should know its own. `[Q]` *Test: the number is written down and dated.*

S-8 . Measure δ and η before promising anything: δ from decay under zero input, η from attributable durable output per unit input. `[δ, η]`

S-9 . Justify the order proxy independently of the input. Per Paper E §6.3 and FVP 46-C/FVP-5, the `Q` proxy must not restate interaction volume. "We talk constantly" is an input measure; "a documented method a third party executed successfully" is an order measure. `[Q]`

S-10 . Require an artefact per period. Periods producing no durable object are running at $\eta \sim 0$ regardless of how they felt. `[η]` *Test: name the artefact for each of the last six.*

S-11 . Measure λ, the convergence rate. Two clean u measurements, six or more periods apart, with no restoration in between. `λ = 1 − (u_1/u_0)^{1/Δt}`. `[u maintenance]`

S-12 . Schedule restoration. Interval no longer than the measured half-life of u. A restoration is a deliberate importation of a state the pair cannot predict from its own: a third party from another lineage in a working session, a rotation into an unfamiliar problem class, an adversarial review by someone differently trained. **A holiday is not a restoration; it changes C, not u.** [u] *Test: restorations are on the calendar, each naming its source of difference.*

S-13 . Measure both private stores, not only the joint one. Per FVP 46-D/FVP-5, reciprocity is unidentifiable from the joint account. [η, reciprocity]

S-14 . Set a reciprocity threshold and act on it. Recommended trigger: a > 0.6931 (a 2:1 directional imbalance) prompts a review of what flows each way. Below that, note and monitor. [reciprocity]

S-15 . Keep Δ_R bounded away from zero. Replies that restate the sender return no information. Where M cannot be ratio-scaled, demote to the ordinal test — "did the reply contain a position I could not have generated?" — and record it as ordinal. [information]

S-16 . Do not add a third person to fix a dyad problem. One head triples the channels (1 -> 3) and raises C at once; it raises u only if the third is oblique to both. Compute h first. [C, u]

S-17 . Compute headroom before requesting an increment. h = η ln N − δQ − C; no increment may exceed h (Cor 4.2). [h]

S-18 . Recognise that h falls as the pair succeeds. δQ grows with Q, so a long partnership spends more of its budget maintaining what it built. Plan the transition to maintenance rather than being surprised by it. [h]

S-19 . Shed Q deliberately when h goes negative. Retiring a shared method, dropping an accumulated obligation — legitimate and underused moves. [Q, h]

S-20 . Every measurement here is opt-in for both parties. CANON §5.4. [ethics]

The Failure Modes

Each: NAME, description, damaged parameter, signature, measurement, fix.

FM-46.1 . THE MIRROR PAIR

Description. Two people selected for similarity — same training, lineage, incentives — whose agreement is read as corroboration. **Damages.** u −> 0, therefore S −> 0. **Signature.** Pre-discussion estimates agree within noise; neither can name a question on which they expect to differ; meetings are short, pleasant, and produce no revisions. **Measurement.** Copy-rule predictive accuracy above 0.9 on held-out domain questions. **Fix.** Import. The pair has no internal source of u (§46.15). A practitioner from an unrelated trade in a working session, on a schedule (S-12), not as a one-off.

FM-46.2 . THE OPPOSITION PAIR

Description. Systematic mutual contradiction mistaken for rigour. ρ ∼ −1. **Damages.** u −> 0, C up. **Signature.** Each position is predictable by inverting the other's; conflict constant, positions

stable, nothing moved in a year. **Measurement.** Inversion-rule predictive accuracy above 0.9. **Fix.** As FM-46.1 — the same fault with a sign flip — plus cut C: the heat is dissipation and buys nothing.

FM-46.3 . THE CONVERGED PAIR

Description. A pair that began oblique and became parallel through successful collaboration. The commonest failure in long partnerships and the least likely to be diagnosed, because every step that produced it was correct. **Damages.** u, via λ. **Signature.** Q peaks then declines while headroom stays positive; fatigue reported; the decline attributed to market conditions or workload. **Measurement.** Two u measurements six or more periods apart (S-11); a halving is decisive. **Fix.** Restoration at or below the measured half-life (S-12). Per §46.17, the first restoration recovers most of what is available.

FM-46.4 . THE ONE-WAY CHANNEL

Description. Conversion flows overwhelmingly one way; one private store grows, the other reaches a negative equilibrium. **Damages.** η directionally; Q for one member. **Signature.** The joint account looks healthy. One partner reports growth, the other depletion, and both reports are accurate. **Measurement.** Separate private stores; a = |ln(η_AB/η_BA)| (S-13). **Not detectable from the joint account — FVP 46-D/FVP-5. Fix.** Reallocate what flows each way. Where the asymmetry is structural — a genuine mentorship — say so and set an end date rather than running it as a partnership.

FM-46.5 . THE PLEASANT VOID

Description. High-u, high-S, η ~ 0. Genuinely different people who enjoy each other and build nothing. **Damages.** η. **Signature.** "We had a great conversation." No artefact for six periods; Q flat or declining despite strong measured u. **Measurement.** Artefact count (S-10); η = ΔQ_attributable / S_input. **Fix.** Chapter 4's capture mechanisms at pair scale: decision log, written summary, changed default, built object. **Nearly always the cheapest intervention available, and it should be tried before anything is done to u.**

FM-46.6 . THE ECHO PARTNER

Description. One member returns the other's own signal. Δ_R ~ 0. **Damages.** information return; downstream, u. **Signature.** Every reply is a restatement, elaboration or endorsement; nothing comes back that the sender could not have generated. **Measurement.** Δ_R where M is ratio-scaled; the ordinal demotion of S-15 otherwise. **Fix.** Structural, not personal: independent-estimate-before-discussion (Ch 22) at pair scale — both write first, then compare. Defeats the echo mechanically, with no confrontation.

FM-46.7 . THE SUB-ANGLE PAIR

Description. A pair below its own u_min: the steady state is negative and the store erodes at rate δ whatever anyone does. **Damages.** Q, terminally. **Signature.** Everything built two years ago is degrading; effort produces maintenance, never gain. **Measurement.** u_min = (e^{C/η}

– 1)/(N–1) versus measured u (S-7). **Fix.** Raise u (import), raise η (capture), or cut C. Raising N is logarithmic and the worst buy (Ch 8). **If none is available, Corollary 3.4 applies: no arrangement of these two people prevents decline.**

FM-46.8 . THE THIRD-PERSON REFLEX

Description. Adding a third member to fix a dyadic deficit. Channels go 1 -> 3, C rises at once, u rises only if the newcomer is oblique to both. **Damages.** C, and often u, if the third is a lineage-mate of either. **Signature.** Headroom falls after the addition; output does not rise. **Measurement.** h before and after (S-16). **Fix.** Prevention: measure the newcomer's u against **both** members first. A third oblique to one and parallel to the other buys less than the channel count costs.

FM-46.9 . THE UNJUSTIFIED PROXY

Description. The order measure repackages the input measure — "we have great shorthand" counted as Q, when shorthand is downstream of interaction volume. **Damages.** measurement validity; η unidentifiable. **Signature.** η near 1 with little residual; excellent training fit, held-out failure. **Measurement.** Paper E §6.3 check: is Y derivable from U without anything outside the pair? **Fix.** An order proxy with an external referent: work a third party can execute, a decision whose outcome is later observable, a physical artefact.

FM-46.10 . THE UNBADGED IDENTIFICATION

Description. Treating the Dyadic Identification (§46.2) as though it were Theorem 4.3(a). **Damages.** the argument's warrant. **Signature.** "The mathematics proves this partnership cannot work." **Measurement.** Read the sentence: does it assert something about vectors or about people? **Fix.** Restate with the badge on. The theorem is exact about states; the identification of people with states is ours, is testable, and has not passed FVP-7.

Worked Example — A Two-Person Design Practice

A constructed case, carried end to end. **Every number below is the example's own; none is drawn from a real firm and none should be cited as empirical.** The arithmetic is reproducible from the stated parameters.

The system

Two people found a small design practice — one trained as a structural engineer, one as an industrial designer — in a domain they define as N = 12 categories: structure, material, cost, fabrication, use, maintenance, regulation, client, schedule, aesthetics, environment, disposal. They review quarterly.

1. Encoding (Paper E §5.1's five required choices)

Choice	Setting
State encoder	Each partner's weighting across the 12 domain categories on a project brief, elicited independently before discussion
Interaction observable	Pre-discussion disagreement across the 12 categories, per project
Order observable `Q`	Durable practice assets: documented details a third party can execute, standing supplier relations, reusable specifications, precedents that survive to the next project
Time scale	One quarter = one period
Cost observable `C`	Coordination time, rework, the overhead of joint decision-making, and the maintenance of the asset library

`Q` is measured by third-party executability, external to the pair's own interaction volume, which satisfies S-9 and Paper E §6.3.

2. Baseline measurement, quarter 0

```
N     = 12       ln N = 2.4849
u     = 0.62     measured by pre-discussion disagreement over 14 briefs
eta   = 0.55     order-units per nat
delta = 0.12     per quarter
C     = 0.34     order-units per quarter
Q_0   = 0.00     new practice
```

Geometry:

```
N_eff = 1 + 11(0.62) = 7.8200
S_0   = ln(7.8200)   = 2.0567 nats
```

Two people with an effective count of 7.82. In degrees, `u = 0.62` is `θ = 51.9` — well inside the productive band of §46.6.

Decomposition (S-6): roughly `u_cross = 0.50` (two trades) and `u_div = 0.12` (each defers to the other inside the other's specialism). Checking Theorem 4.1 at `N = 12`:

```
f(0.50) = 1.8718     f(0.12) = 0.8416     f(0.62) = 2.0567
max = 1.8718  <=  2.0567  <=  2.7134 = sum      [Thm 4.1 satisfied]
```

The pair is rich in `u_cross` and thin in `u_div`. Flag it now; it becomes the story.

3. The minimum viable angle (S-7)

```
S_min = C / eta = 0.34 / 0.55 = 0.6182
u_min = (e^0.6182 - 1) / 11   = 0.0778        theta_min = 16.2 degrees
N_eff at u_min = 1.8556
```

The practice's survival line is 16.2 degrees. At quarter 0 they are at 51.9 degrees, comfortably clear.

4. Convergence measured (S-11)

Two clean u measurements, quarters 0 and 8, no restoration in between:

```
u_0 = 0.62    u_8 = 0.378    lambda = 1 − (0.378/0.62)^{1/8} = 0.0600 per quarter
half−life of u  = ln 2 / (−ln(1−0.06)) = 11.2 quarters
```

Their difference has a half-life of under three years. They did nothing wrong; they learned each other's trades, which is what a good cross-disciplinary pair does.

5. The trajectory with no restoration

Running `Q_{t+1} = (1−0.12)Q_t + 0.55 S_t − 0.34` with `u_t = 0.62(0.94)^t`:

t	u	N_eff	S	Q	h	step
0	0.6200	7.820	2.0567	0.0000	+1.0267	+0.7912
2	0.5478	7.026	1.9496	1.4579	+0.8518	+0.5574
4	0.4841	6.325	1.8445	2.4766	+0.7295	+0.3773
8	0.3779	5.157	1.6404	3.5785	+0.5973	+0.1328
10	0.3339	4.673	1.5419	3.8009	+0.5706	+0.0519
12	0.2951	4.246	1.4459	3.8719	+0.5621	−0.0094
16	0.2304	3.534	1.2625	3.6993	+0.5828	−0.0896
20	0.1799	2.979	1.0914	3.2684	+0.6345	−0.1319
24	0.1404	2.545	0.9340	2.7064	+0.7019	−0.1511

6. Reading the trajectory

The practice peaks at quarter 12 with `Q = 3.87` and declines thereafter, holding `Q = 2.71` by quarter 24 — roughly its quarter-5 position, six years in. Now look at the headroom column, which is the diagnostic point of the whole example:

> **Headroom is positive at every single step, and it is HIGHER at quarter 24 (+0.70) than at quarter 10 (+0.57).** The master gauge of Volume One never once signalled a problem. It could not: `h = η ln N − δQ − C` contains no term for `u`. Headroom measures the *budget*, not the *input*. A pair can have ample budget and nothing to spend it on.

Headroom tells you what you could bank; `u` tells you what there is to bank. A partnership dies with a healthy balance sheet.

Headroom *improves* as the practice declines, because `δQ` shrinks with `Q`: the gauge gets more encouraging exactly as the situation gets worse. Not a defect in the gauge — a demonstration that one gauge is not a dashboard.

7. What the partners will say at quarter 20

They will report that the work has become repetitive, that clients are less interesting, that the market has changed, and that they are tired. All four reports will be sincere. None will mention that their effective count has fallen from 7.82 to 2.98. **u is the parameter that decays without any experience of decaying**, because from the inside convergence feels like competence.

8. The intervention

Apply S-12: restore u to 0.62 on a schedule at or below the measured 11.2-quarter half-life. Running the same model with restoration every k quarters, Q at quarter 24:

restoration interval	Q at t=24	gain vs never
never	2.7064	—
every 12 quarters	4.7070	+2.0006
every 8 quarters	5.3283	+2.6219
every 6 quarters	5.6209	+2.9145
every 4 quarters	5.8991	+3.1927
every 2 quarters	6.1615	+3.4551

Restoring once every twelve quarters — three years — captures 58 per cent of the total available gain; a sixfold increase in effort, to two-quarterly, captures the remaining 42 per cent. Concavity again: do it at all before arguing about how often.

Concretely, a restoration here is a fabricator, a regulator, a maintenance contractor or a practitioner from an unrelated trade, brought into a live working session with standing to disagree, for one project. Not a lecture. Not a conference. Someone whose position on the twelve categories neither partner can predict, inside the actual work.

This addresses the deficit flagged in step 2: their `u_cross` was strong and is eroding as they learn each other's trades, while `u_div` was thin from the start. **A third trade rebuilds `u_cross`; a dissenting practitioner from one of their own trades rebuilds `u_div`. They need both, and the diagnosis says which is more urgent.**

9. What would falsify this

The prediction is specific and preregisterable:

1. u measured at quarter 8 and quarter 16, with restoration events at quarters 4 and 12, will be higher than the no-restoration decay curve predicts for those quarters (0.378 and 0.230).
2. Q measured by third-party executability at quarter 24 will exceed the quarter-12 peak of 3.87 under a restoration schedule of 12 quarters or shorter.
3. The fitted model $Y_{t+1} = aY_t + bF(U_t) - \hat{C}_t + \varepsilon_t$, frozen after quarter 12, will beat all four baselines on quarters 13-24 held out.

The framework is falsified for this practice if the shuffled-input control — the same `F(U_t)` series with its time order randomised — predicts held-out Q as well as the fitted model. That would mean the timing of the pair's non-parallelism carried no information about the timing of what it built, which is precisely the claim the storage law makes.

Be explicit about what has and has not been established. The arithmetic is exact; the parameters are constructed. **FVP-7 has not been run on any of it, at dyad scope, by anyone**, and until it has, this chapter is a specification for how to find out rather than a report of what is known. That distinction is the book's warrant and it is not negotiable.

The dyad is the atom: two people, one channel, one angle, and a theorem saying the angle is the whole story. **The wedding of like with like is a non-event, and it is a non-event whether the likeness is agreement or its mirror.** Chapter 47 takes Volume Two's other equation — the moment equation — and asks what it measures when the perceiver is an organisation.

Chapter 47 . The Group Moment

Volume Two rests on two equations. Chapter 34 gave the second of them:

```
M = (D / P) . E
```

and gave it at the scale of one person: `D` the frame a mind brings, `P` the intake that mind achieves, `E` the environment, `M` the moment as lived. This chapter takes the same equation, keeps every symbol, changes only the referent of each, and asks whether it still says anything.

It does. `D` becomes **shared doctrine** — the standing model an organisation carries about the world it operates in. `P` becomes **collective perception**: the intelligence the organisation actually gathers, audited for redundancy. `E` is the environment the organisation is in. `M` is **the situation as the group constructs it** — the thing the organisation acts on, which is never the environment itself.

This chapter is the proof that the same equation operates at both scopes, and the proof is not rhetorical: the group version acquires a term the personal version does not have, and that term is the surplus `F` from Paper A. Doctrine over data, at group scale, contains the interaction geometry of the group inside it. That is the join between Volume Two's two equations, and it is computed here.

The Law

47.1 The equation, restated at group scope

```
   M = ( D / P ) . E

  D  shared doctrine      the standing model: what the organisation already asserts
   P  collective perception  intelligence gathered, audited for redundancy
   E  environment            the world the organisation is in
   M  the constructed        the situation the organisation acts on
      situation

   r  = D / P                the doctrine-to-data ratio (dimensionless)
  Δ  = | ln r |          distortion, in nats, same register as everything else

   ln M = ln D − ln P + ln E    the log form (CANON_INTERNAL §1, register 4)
```

Status badge. `M = (D/P).E` is **[POSTULATE — the author's, stated in the open]**. It is not a theorem of Papers A-E and nothing in this manual derives it. It is a proposed structure whose consequences are computable and, by the protocol of §47.11, falsifiable. Everything downstream in this chapter that is *proven* comes from Paper A or Paper E and is labelled as such; everything else inherits the postulate's status.

47.2 The commensurate-units convention

FVP-1 requires `D` and `P` to share units so that `r` is dimensionless. At personal scope Chapter 34 handles this by magnitude estimation. At group scope there is a cleaner route, because organi-

sations already count the thing.

> **THE COMMENSURATE-UNITS CONVENTION [POSTULATE].** Express `D` and `P` in the same currency: **effective independent evidence-channels per decision period.**
>
> - `P` is what the organisation takes in: the number of channels it consults, audited for redundancy, times the mean intake per channel.
> - `D` is what the organisation's standing model asserts about the same decision space *without looking*: the decision-relevant propositions the doctrine fixes in advance, each weighted by the fraction of decisions it constrains, summed. This is how many channels' worth of constraint the doctrine supplies for free.
> - `E` is the decision-relevant structure actually present in the environment, in the same currency.
>
> All three are then counts with a true zero, hence ratio-scaled (FVP-2), and `M` inherits `E`'s units as CANON_INTERNAL §4 requires.

This convention is doing real work and it is the weakest joint in the chapter. It is badged, and §47.10's FVP-5 finding shows exactly where it can fail.

47.3 The aggregation postulate — where Paper A enters

`P` is not headcount. A hundred analysts reading the same wire service are one channel. This is Chapter 7's effective count, and it is the entire reason the surplus law belongs in a chapter about perception.

> **THE AGGREGATION POSTULATE [POSTULATE].** Collective perception is the mean per-channel intake times the effective count of the channels:
>
> ```
> P = p_bar . N_eff      N_eff = 1 + (N−1) u      (Paper A, Thm 2.1) [PROVEN form]
> ```
>
> where `N` is the number of distinct domains the organisation's intelligence can express, `u` the mean non-parallelism between its channels, and `p_bar` the mean intake per genuinely independent channel.

Taking logs and substituting:

```
ln P = ln p_bar + ln N_eff = ln p_bar + F       since F = ln N_eff exactly

ln r = ln D − ln p_bar − F                          THE JOIN
```

The surplus of the group appears in the doctrine-to-data ratio as a subtractive term, one nat for one nat. Every nat of interaction surplus the organisation's channels generate reduces its log-distortion by exactly one nat. That is not an analogy between the two equations; it is an identity, and it holds because `F` is defined as `ln N_eff` and `ln` is the register both formulas live in (CANON_INTERNAL §1).

Two consequences follow immediately, and both are hard.

Consequence 1 — the correction ceiling. Since `0 <= F <= ln N` (Paper D, Cor 2.2), the total reduction in `ln r` available from interaction geometry alone is at most `ln N`:

```
N =  4  ->  max reduction 1.3863 nats   -> r divided by at most  4
N =  8  ->  max reduction 2.0794 nats   -> r divided by at most  8
```

```
N = 12  ->  max reduction 2.4849 nats   -> r divided by at most 12
N = 24  ->  max reduction 3.1781 nats   -> r divided by at most 24
N = 40  ->  max reduction 3.6889 nats   -> r divided by at most 40
```

Consequence 2 — the infeasibility condition. Setting `r = 1` requires `N_eff = D/p_bar`, and `N_eff <= N`. Therefore:

> **If `D / p_bar > N`, the organisation cannot be calibrated by any arrangement of its people.** The residual distortion is at least `ln(D/p_bar) − ln N`, and no recruitment, rotation, red team or restructuring closes it. The only remaining move is to reduce `D`.

This is the exact analogue, in the perceptual register, of Paper E Corollary 3.4's statement that if `C_t > η ln N − δQ_t` no interaction geometry prevents decline. **Both are `ln N` ceilings, and both say the same thing: geometry is bounded, and past the bound you must change what you are carrying, not who is carrying it.**

47.4 The contact postulate — how the moment feeds the storage law

CANON_INTERNAL §8 requires the moment to serve as the storage law's input: `Q_{t+1} = (1−δ)Q_t + ηM_t − C_t`, with `M` normalised to `[0, ln N]`. The normalisation must be stated, not assumed.

> **THE CONTACT POSTULATE [POSTULATE].** The surplus a group can bank in a period is the ceiling `ln N` attenuated by its distortion:
>
> `S_t = ln N . e^{−Δ_t}` `Δ_t = | ln r_t |`

That exponential has an exact and much more legible form. Since `e^{−|ln r|} = min(r, 1/r)` identically — verified at `r = 0.1, 0.25, 0.5, 0.8, 1.0, 1.25, 2.0, 4.0, 10.0` to machine precision:

```
S_t / ln N  =  min(D, P) / max(D, P)
```

A group banks at the rate of the scarcer of frame and intake. That is Liebig's law of the minimum, arrived at from the log-ratio rather than assumed: doctrine you cannot feed with evidence and evidence you cannot frame are equally useless, and the binding one is whichever is smaller. The attenuation is symmetric, which is the postulate's substantive commitment and the thing §47.11's falsification test attacks first.

The band table, Δ to attenuation:

Δ	r (either side)	S / ln N
0.00	1.000	1.0000
0.05	1.051 or 0.951	0.9512
0.10	1.105 or 0.905	0.9048
0.20	1.221 or 0.819	0.8187
0.35	1.419 or 0.705	0.7047
0.50	1.649 or 0.607	0.6065
0.69	1.994 or 0.502	0.5016
1.00	2.718 or 0.368	0.3679
1.39	4.015 or 0.249	0.2491
2.00	7.389 or 0.135	0.1353

A two-to-one doctrine-to-data ratio halves what the organisation can bank. That is the number to carry out of this chapter.

47.5 The honesty clauses at group scope

1. **`M = (D/P).E` is a postulate, not a theorem.** Papers A-E prove nothing about it.
2. **The Aggregation Postulate is a postulate.** That P is linear in `N_eff` is chosen for the same reason S3 was chosen in Paper A: it is the simplest assumption that makes an aggregate behave like a count. If measurement replaces it, everything downstream is rebuilt.
3. **The Contact Postulate is a postulate**, and its symmetry — over-framing and under-framing attenuating equally at equal `|ln r|` — is its most testable and most doubtful commitment.
4. **`N_eff = 1 + (N−1)u` is affine because S3 says so.** Design axiom, chosen.
5. **The surplus functional is a lower bound** on the ensemble's information. P computed this way is a floor on the organisation's intake, not its ceiling.
6. **Surplus is not Shannon entropy or mutual information** (Paper D). Do not read P as channel capacity.
7. **No thermodynamic claim is made anywhere in this chapter** (Paper E honesty clause).
8. The predictive-processing correspondence noted in CANON_INTERNAL §4 — precision-weighted priors against sensory evidence — is a **structural correspondence between formalisms, not an identity**, and predictive processing is an active and contested research programme. It is offered as a striking parallel and nothing more.

The Engineering

47.6 The four limiting cases, worked at group scale

FVP-3 requires every limit to correspond to a real, nameable phenomenon. All four do, and at group scope they are more recognisable than at personal scope, because organisations fail publicly.

CASE 1 — `P -> 0`: doctrine with no intelligence. `M -> inf`.

```
D = 1.0, E = 1.0

P        r          M        Δ
-----    -------    ------   ------
1.000      1.00      1.00    0.0000
0.500      2.00      2.00    0.6931
0.200      5.00      5.00    1.6094
0.100     10.00     10.00    2.3026
0.050     20.00     20.00    2.9957
0.020     50.00     50.00    3.9120
0.010    100.00    100.00    4.6052
0.001   1000.00   1000.00    6.9078
```

The divergence is real and the phenomenon has a name in every language: **an organisation with a model and no collection is not blind, it is certain.** The constructed situation grows without

bound while the environment stays where it is. Behaviourally this reads as institutional dogma, prophecy, and the pathology Chapter 48 calls over-framing.

Note what the arithmetic says about the *feel* of the condition. `M` is large. The organisation experiences a rich, detailed, confident picture. **Nothing in the subjective quality of `M` distinguishes a well-informed organisation from a hallucinating one; only the ratio does.** This is the single most important reason the ratio has to be measured rather than sensed.

Note also that `P -> 0` does not require anyone to stop collecting. Under the Aggregation Postulate, `P = p_bar . N_eff`, and `N_eff -> 1` as `u -> 0`. **An organisation whose channels become parallel has `P` collapse to a single channel's worth while its collection budget stays flat.** That is Chapter 48's mechanism and it is stated here because it belongs to the limit, not to the pathology.

CASE 2 — `D -> 0`: intelligence with no frame. `M -> 0`.

```
P = 1.0, E = 1.0

D        r         M        Δ
-----    ------    ------   ------
1.000    1.0000    1.0000   0.0000
0.500    0.5000    0.5000   0.6931
0.200    0.2000    0.2000   1.6094
0.100    0.1000    0.1000   2.3026
0.050    0.0500    0.0500   2.9957
0.010    0.0100    0.0100   4.6052
0.001    0.0010    0.0010   6.9078
```

`M -> 0` and the organisation constructs no situation at all. This is the phenomenon CANON_INTERNAL §4 names precisely: **sensation is not experience.** At group scope: an organisation holding an enormous quantity of data and no model of what it is for produces no decision. The data lake with no question. The dashboard nobody can act on. The agency that collects everything and concludes nothing.

The limit is genuinely `0` and not merely small, and that is worth stating: a group with no doctrine does not act badly, **it does not act**. The paralysis of the perfectly open-minded institution is a prediction of the formula, not a moral observation.

CASE 3 — `E -> 0`: no environment. `M -> 0` for any `r`.

The organisation's environment has gone. Its market has vanished, its constituency has dissolved, its problem has been solved or has migrated elsewhere. **`M = 0` regardless of how good its doctrine and intelligence are**, because a ratio multiplied by nothing is nothing.

This limit is often mistaken for the others, and the distinction is diagnostic. An organisation with `E -> 0` has excellent internal metrics and produces nothing that lands. Its doctrine may be perfectly calibrated to its intake — `r = 1`, distortion zero — and it will still generate no moment. **Calibration is not sufficiency. A perfectly calibrated instrument pointed at nothing reads nothing.**

CASE 4 — D = P: M = E. The fixed point.

```
r = 1  ->  Δ = 0  ->  M = E  ->  S = ln N
```

When doctrine matches intelligence, the situation the organisation constructs equals the environment it is in. That is undistorted contact, stated as an equation, and CANON_INTERNAL §4 is right to call it the formula's most important content.

There is a structural result at this fixed point that does not appear at personal scope, and it is the most surprising thing in the chapter. Put `r = 1` into the Contact Postulate: `S = ln N`, the maximum. Then the storage law's steady state is `Q_inf = (η ln N − c)/δ`, and headroom at that steady state is:

```
h = η ln N − δ Q_inf − C
  = η ln N − δ (η ln N − C)/δ − C
  = η ln N − η ln N + C − C
  = 0        exactly
```

> **A perfectly calibrated organisation converges to exactly zero headroom.** Not by coincidence and not by mismanagement: by arithmetic. Perfect contact drives the input to the ceiling, the ceiling drives the store to its maximum, and the maximum store consumes the entire budget in maintenance.

This is Volume One's Chapter 6 thesis — success consumes its own headroom — arriving from the perceptual side. It also warns against a naive reading of `r = 1` as an unqualified goal. Calibration maximises what you bank; it does not leave you room to build something new with it. **An organisation that wants headroom at equilibrium must run below its maximum store, which means either shedding Q deliberately (Ch 23) or accepting r != 1.** The manual states the trade rather than pretending it away.

47.7 The calibration requirement, computed

Solve `r = 1` for the non-parallelism required. `N_eff = D/p_bar`, so:

```
u* = ( D/p_bar − 1 ) / ( N − 1 )        feasible only if D/p_bar <= N
```

The full table, across domain sizes and doctrine loads:

D/p_bar	N=4	N=8	N=12	N=16	N=24	N=40
1.5	0.1667	0.0714	0.0455	0.0333	0.0217	0.0128
2.0	0.3333	0.1429	0.0909	0.0667	0.0435	0.0256
3.0	0.6667	0.2857	0.1818	0.1333	0.0870	0.0513
4.0	1.0000	0.4286	0.2727	0.2000	0.1304	0.0769
6.0	--	0.7143	0.4545	0.3333	0.2174	0.1282
8.0	--	1.0000	0.6364	0.4667	0.3043	0.1795
12.0	--	--	1.0000	0.7333	0.4783	0.2821
16.0	--	--	--	1.0000	0.6522	0.3846
24.0	--	--	--	--	1.0000	0.5897
40.0	--	--	--	--	--	1.0000

`"--"` = INFEASIBLE: `D/p_bar` exceeds `N`; no geometry calibrates this organisation.

Three readings of that table, each an engineering rule.

Reading one: doctrine is expensive in u. An organisation whose standing model supplies six channels' worth of constraint, working an N = 12 domain, needs u = 0.4545 — mean pairwise non-parallelism of 0.45 across its intelligence channels, θ = 42.4 degrees. That is a high bar. Most organisations measured by the protocols of Chapter 9 come in an order of magnitude below it.

Reading two: a large N is a doctrine subsidy. The same D/p_bar = 6.0 needs u = 0.4545 at N = 12 and only u = 0.1282 at N = 40. Breadth of domain buys tolerance for doctrine. But note Chapter 8's warning: N moves the ceiling only logarithmically and is the most expensive lever there is. **Raising N from 12 to 40 to relieve a doctrine problem is a real move and a bad buy.**

Reading three: the infeasible region is large and it is where heavy institutions live. A long-established body carrying D/p_bar = 16 in an N = 12 domain is in the "–" zone. Its residual distortion floor is:

```
Δ_min = ln(D/p_bar) − ln N

N=12, D/p_bar=18  ->  Δ_min = 0.4055        S <= 0.6667 ln N
N=12, D/p_bar=24  ->  Δ_min = 0.6931        S <= 0.5000 ln N
N=12, D/p_bar=48  ->  Δ_min = 1.3863        S <= 0.2500 ln N
N=24, D/p_bar=48  ->  Δ_min = 0.6931        S <= 0.5000 ln N
N=24, D/p_bar=96  ->  Δ_min = 1.3863        S <= 0.2500 ln N
```

An institution carrying four times its domain size in doctrine cannot bank more than a quarter of its ceiling, no matter who it hires. The only lever left is D, and Chapter 48 specifies how to pull it.

47.8 Worked mechanics of the join

Fix D = 6.0, p_bar = 1.0, N = 12, and let the group's own geometry vary. This is the join running end to end: u sets F, F sets ln r, ln r sets Δ, Δ sets S.

u	F	ln r	r	Δ	M/E
0.0000	0.0000	+1.7918	6.0000	1.7918	6.0000
0.0200	0.1989	+1.5929	4.9180	1.5929	4.9180
0.0500	0.4383	+1.3535	3.8710	1.3535	3.8710
0.1000	0.7419	+1.0498	2.8571	1.0498	2.8571
0.2000	1.1632	+0.6286	1.8750	0.6286	1.8750
0.3000	1.4586	+0.3331	1.3953	0.3331	1.3953
0.4167	1.7198	+0.0720	1.0746	0.0720	1.0746
0.4545	1.7918	0.0000	1.0000	0.0000	1.0000
0.5000	1.8718	−0.0800	0.9231	0.0800	0.9231
0.7000	2.1633	−0.3716	0.6897	0.3716	0.6897
0.9000	2.3888	−0.5970	0.5505	0.5970	0.5505
1.0000	2.4849	−0.6931	0.5000	0.6931	0.5000

Read the column of ln r against the column of F. **They move one for one, in opposite directions, exactly.** ln r = ln 6 − F = 1.7918 − F on every row. The surplus of the group *is* the

de-framing of the group, in the same units, with no conversion constant.

Read the last four rows for the other half of the story. **Past u = 0.4545 the organisation over-corrects into under-framing.** At u = 1.0 — perfectly orthogonal channels — the same doctrine load produces r = 0.5, a two-to-one intake surplus, and M/E = 0.5: the organisation now under-constructs the situation. It is drowning.

That is a genuinely useful and non-obvious result. **Diversity is not monotonically good in this model.** It reduces ln r, and reducing ln r is good only while r > 1. An organisation already under-framed and adding channels is making its condition worse, and the formula says so plainly. This is the sense in which Chapter 48's two pathologies require *opposite* corrections.

47.9 The moment as storage input, computed

Take a group with η = 0.42 order-units per nat, δ = 0.14 per period, C = 0.55 order-units per period, N = 12 so ln N = 2.4849. The minimum input for a non-negative equilibrium is S_min = C/η = 1.3095 nats.

S	ηS	Q_inf
0.0000	0.0000	−3.9286
0.3000	0.1260	−3.0286
0.6000	0.2520	−2.1286
0.9000	0.3780	−1.2286
1.2000	0.5040	−0.3286
1.5000	0.6300	+0.5714
1.8000	0.7560	+1.4714
2.1000	0.8820	+2.3714
2.4849	1.0437	+3.5261

Now convert distortion to input by the Contact Postulate, S = ln N . min(r, 1/r), and read off the distortion at which this group stops being able to hold anything at all:

```
S_min / ln N = 1.3095 / 2.4849 = 0.5270
min(r, 1/r) >= 0.5270   ->   0.527 <= r <= 1.898   ->   Δ <= 0.6406
```

> **This organisation must keep its doctrine-to-data ratio between 0.53 and 1.90. Outside that band its steady-state store is negative and its institution erodes at rate δ regardless of effort.**

That is a concrete, computable operating envelope derived from three measured parameters and a badged postulate. It is the kind of statement the manual exists to produce.

47.9.1 The two-track ledger at group scope, and what it does to the target

The envelope above is stated on Q, and Q is the wrong ledger to optimise. Chapter 34 §34.21 establishes the canon — **over-framing INCREASES stored order** — and Chapter 40 §40.2 turns it into an instrument. Both transfer to group scope without modification, because the storage law does not know what scale it is running at:

```
       ALL  STORED  ORDER          Q_{t+1}     =  (1−δ)  Q_t     +  η  M_t  −
C_t          M_t = min(r_t E_t, ln N)
```

```
       VERIDICAL  ORDER            Q^V_{t+1}  =  (1−δ)  Q^V_t  +  η  V_t  −
C_t          V_t = E_t . min(r_t, 1)

       STORED  ERROR                 Z_{t+1}     =  (1−δ)  Z_t  +  η  (M_t  −
V_t)          [the cost term cancels]
    CORRECTION LOAD       C_t       = c_0 + κ Z_t
```

An over-framed institution accumulates MORE stored order than a calibrated one. It is simply not true stored order. The doctrine, the procedure manual, the settled interpretation, the confident forecast — all of it banks at rate η whether or not the environment supplied it, because there is no term in the recursion that asks. **Dogma banks faster than knowledge, and the bill arrives on a lag, as correction load rising with accumulated error.** That lag is why an over-framed body looks excellent on its own metrics for years, and it is the mechanism behind the doctrine-to-data pathology Chapter 48 diagnoses.

This sharpens §47.6's result into something stronger and less comfortable. §47.6 showed that `r = 1` is not an unqualified goal because it converges to zero headroom. The two-track ledger shows that a body optimising raw `Q` is not merely failing to aim at `r = 1` — **it is being actively driven away from it.** Above the fixed point `V` is flat at `E` while `M` keeps climbing, so the veridical track stops improving at `r = 1` and the raw track does not. Differentiating the steady state for `r > 1`:

```
   Q_inf = (η M − c_0 − κ η (M − V)/δ) / δ       −>       dQ_inf/dM = η (1 −
κ/δ) / δ
```

The knife edge is exactly κ = δ. Where correction load accrues more slowly than stored order decays, raw `Q_inf` rises without bound in `r`, and every incentive attached to raw `Q` — headcount, published procedure, doctrinal coverage, confident forecasts issued — is an incentive to over-frame. That is not a claim about anyone's character. It is the arithmetic of what the organisation is being paid in.

> **REQUIREMENT. Report `Z` alongside `Q`, or report `Q^V` instead of `Q`. A body that reports only its total store is reporting a number that rises fastest when it is most wrong.** Chapter 40 §40.2 gives the instrument at personal scope and it transfers unchanged; Chapter 48 gives the correction protocols; Chapter 50 §50.12 records that the two-track requirement holds at both scopes.

47.10 The formulas of this chapter, with their FVP tables

One verdict governs three of the five tables below and is stated before them rather than buried inside them. Chapters 52 and 53, running the Formula Validation Protocol on the Moment Equation, return **DEMOTE on gate 2**, and the verdict carries to group scope. This chapter argued in §47.2 that channel-counting gives `D` and `P` a true zero, and that argument is here **corrected**: a count supplies a zero and a unit *within one declared encoder*, but nothing links two encoders, and without linkage the magnitude of `r` moves when the encoder moves. An organisation's channel count is more auditable than a person's frame, which is a real advantage at FVP-5; it is not the same thing as a scale.

> **NOT REPORTABLE AS MEASURED QUANTITIES, at group scope as at personal scope:** the magnitude of `D/P`; the distortion `Δ = |ln r|` as a number; and the quantitative form of the join from the moment into the storage law. Every figure in

§47.7, §47.8 and §47.9 — the u∗ table, the Δ_min floors, the 0.53–1.90 envelope — is an **illustrative construction demonstrating the structure of the model**, and is to be read as such.

STILL IN FORCE: the four limiting cases of §47.6; the D = P fixed point and its zero-headroom consequence; the ordering over-framed / calibrated / under-framed; the *sign* of ln r; and the algebraic identity ln r = ln D − ln p_bar − F, which is an identity and does not depend on any scale claim. **What is demoted is magnitude, not structure.**

The route out is Chapter 54, Theorem 54.3, given badged Postulate SO1: measure D and P as Rasch measures on one common linked scale and Δ = |ln(D/P)| = |θ_D − θ_P| becomes genuinely ratio-scaled, because Rasch separation supplies the true zero division requires. §54.12 prices it at **approximately 123 items per term for a half-logit distortion at 95 per cent.** At institutional scope the items are decisions with later-observable answers, which is the one respect in which the group version is genuinely cheaper: **an organisation with a decision log is already accumulating the item bank.** That is the specific study which would convert this chapter's envelope from an illustration into a measurement.

FVP TABLE 47-A — The group moment M = (D/P).E

Gate	Finding
FVP-1 DIMENSIONS	Under the Commensurate-Units Convention (§47.2), D, P, E are all in effective evidence-channels per period. D/P dimensionless; M inherits E's units. Both sides agree. **PASS, and it passes ONLY because the convention was declared. Without it the formula is dimensionally void.**
FVP-2 SCALE TYPE	Division requires ratio scale with a true zero on a **common** scale. Channel-counts supply a zero and a unit within one declared encoder — zero channels, zero doctrine, zero environment are all meaningful — but supply no linkage between encoders, so the magnitude of r is encoder-dependent. **DEMOTE (Ch 52, Ch 53): the magnitude of D/P and of Δ is withdrawn as a measured quantity; the ordering, the limits, the fixed point and the sign of ln r stand.** Where an implementation substitutes Likert-style ratings for counts the position is worse still — division is impermissible outright. **Route out: Ch 54 Thm 54.3, Rasch logit difference, ~123 items per term for half a logit, with decisions as items.**

Gate	Finding
FVP-3 LIMITING CASES	All four worked in §47.6. `P -> 0`: `M -> inf`, the certain organisation. `D -> 0`: `M -> 0`, the paralysed one. `E -> 0`: `M -> 0`, the obsolete one. `D = P`: `M = E`. Every divergence names a real, observable institutional condition. **PASS.**
FVP-4 FIXED POINTS	`r = 1 => M = E`, undistorted contact, unique. Plus the derived result of §47.6: at that fixed point, `h = 0` exactly. **PASS, with the trade stated.**
FVP-5 IDENTIFIABILITY	**THIS GATE FAILS AS THE FORMULA IS NAIVELY WRITTEN, and the failure must be printed.** From observations of `M` and `E` alone, only the ratio `r` is recoverable — never `D` and `P` separately. Demonstration: `(D,P) = (6.0, 3.0)`, `(2.0, 1.0)`, `(0.5, 0.25)` and `(120.0, 60.0)` all give `r = 2.0000` and, at `E = 1`, identical `M = 2.0000`. **Four distinct parameter sets, one observation.** *The repair:* an independent anchor on `P` — a direct count of audited channels and their intake — makes `D` recoverable as `D = r . P`. **CONDITIONAL PASS, and only with the anchor. Any deployment that reports `D` without independently measuring `P` is reporting an unidentifiable quantity.**
FVP-6 FALSIFICATION	Paper E §5.3 template. Preregister the encoder for `D`, `P`, `E`; fit on training periods; FREEZE; report held-out. Four baselines: persistence, direct-input (`E` alone), matched autoregressive, and **shuffled-input control** (the `r` series time-shuffled). **Kill conditions: (a) `E` alone predicts held-out decision quality as well as `(D/P).E`; (b) the shuffled-`r` control matches the fitted model; (c) the sign of the residual is independent of the sign of `ln r`.**
FVP-7 REPRODUCIBILITY	The cross-institution test: same operationalisation, different organisations, different analysts — do `D`, `P` recover, and does the same `r` band predict the same behaviour? **UNTESTED. Nobody has run it. This gate decides whether this chapter is science, and stating that plainly is a requirement of CANON_INTERNAL §6.**

FVP TABLE 47-B — Distortion Δ = |ln r|, r = D/P

Gate	Finding
FVP-1 DIMENSIONS	r dimensionless by 47-A; Δ in nats, commensurate with F and with everything else in the volume's log register. **PASS.**
FVP-2 SCALE TYPE	Δ *would be* a magnitude on a ratio scale, with differences interval-scaled — **and this is corrected here: that property is inherited from r, whose magnitude 47-A demotes.** Pending instrumentation Δ is **ORDINAL**: it ranks bodies and periods by distance from contact and does not license the statement that one distortion is twice another. **DEMOTE.** Becomes ratio-scaled exactly when Δ is read as a Rasch logit difference (Ch 54 Thm 54.3) and by nothing cheaper.
FVP-3 LIMITING CASES	r –> 0: Δ –> inf, total under-framing. r –> inf: Δ –> inf, total over-framing. r = 1: Δ = 0. The two divergences are **distinct pathologies mapped to the same magnitude**, which is the measure's one real weakness: Δ alone cannot tell you which way you are wrong. **PASS ONLY IF THE SIGN OF ln r IS REPORTED ALONGSIDE. Never report Δ bare.**
FVP-4 FIXED POINTS	Δ = 0 <=> r = 1 <=> D = P. Unique. **PASS.**
FVP-5 IDENTIFIABILITY	Δ is recoverable wherever r is, and is in fact *more* robustly recoverable, since it survives any rescaling that multiplies D and P by a common factor. **PASS — and this robustness is why Δ is the reportable quantity and D and P are not.**
FVP-6 FALSIFICATION	Preregister that decision quality declines monotonically in Δ on both sides of 1. **Kill condition: no monotone relation, or a relation present on one side and absent on the other** — the latter would falsify the symmetry of the Contact Postulate specifically, which is a useful and separable result.
FVP-7 REPRODUCIBILITY	**UNTESTED.**

FVP TABLE 47-C — Aggregation P = p_bar . N_eff, N_eff = 1 + (N–1)u

Gate	Finding
FVP-1 DIMENSIONS	p_bar in intake-units per channel per period; N_eff a pure count; P in intake-units per period. Matches D's units by the convention. **PASS.**
FVP-2 SCALE TYPE	N a count; u ratio-scaled per Chapter 46's FVP 46-A, conditional on the channel encoder; p_bar a mean of ratio-scaled intakes. **PASS, conditional on the encoder being declared.**
FVP-3 LIMITING CASES	u –> 0: N_eff –> 1, P –> p_bar — **a hundred parallel channels are one channel**, which is Chapter 7's result and Chapter 11's pathology. u –> 1: P –> p_bar . N, the maximum. p_bar –> 0: P –> 0, Case 1 — an organisation that consults many independent sources and takes nothing from any of them. N –> 1: P = p_bar for all u, correct. **PASS; each limit is a real institutional condition.**
FVP-4 FIXED POINTS	u = 0 is the invariant of the parallel organisation and is **globally attracting under the decay of Chapter 11** — u is maintained, not owned. u = 1 is the invariant of full independence and is not attracting at all. **PASS, with the asymmetry noted: the bad fixed point is the stable one.**
FVP-5 IDENTIFIABILITY	N by enumerating declared intelligence domains; u by error-decorrelation across channels on questions with later-observable answers; p_bar by direct intake tally. All three separately obtainable. **But P alone does not identify (p_bar, N_eff)** — a well-read single channel and a barely-read diverse set can give the same P. **Report all three, never P alone. PASS IF ALL THREE ARE REPORTED.**
FVP-6 FALSIFICATION	Preregister that held-out forecast accuracy tracks N_eff rather than channel headcount or channel budget. **Kill condition: headcount or budget predicts held-out accuracy better than N_eff.** This is a clean, cheap test and it is the one an organisation with a forecasting record can run first.

Gate	Finding
FVP-7 REPRODUCIBILITY	**UNTESTED at institutional scale.** Woolley 2010's collective-intelligence factor is the nearest published evidence that group performance tracks something other than member ability, and it is a laboratory, short-task, small-group result. Supporting context; not confirmation.

FVP TABLE 47-D — Contact `S = ln N . e^{−Δ} = ln N . min(r, 1/r)`

Gate	Finding
FVP-1 DIMENSIONS	`Δ` in nats, `e^{−Δ}` dimensionless in (0,1], `ln N` in nats, `S` in nats. Matches the storage law's input requirement exactly. **PASS.**
FVP-2 SCALE TYPE	Requires `r` ratio-scaled, which 47-A demotes. `S` is in nats and would be ratio-scaled with a true zero, but it is computed from `r` and can be no better than its input. **DEMOTE: the postulate's ORDERING — contact banks more than distortion, in either direction, monotonically — stands on ordinal information alone. The banked-nats figures of §47.9 do not.** Restored by Ch 54 Thm 54.3.
FVP-3 LIMITING CASES	`Δ –> 0`: `S –> ln N`, the ceiling, achieved only at exact calibration. `Δ –> inf`: `S –> 0`, the organisation banks nothing — true in both directions, the certain institution and the flooded one alike. Bounded above by `ln N` as Paper D Cor 2.2 requires, so the postulate is **consistent with the proven range and does not smuggle in surplus above the ceiling. PASS.**
FVP-4 FIXED POINTS	`S = ln N` at `r = 1`, and the derived consequence `h = 0` at the resulting `Q_inf` (§47.6). **PASS.**
FVP-5 IDENTIFIABILITY	`S` is a deterministic function of `r` and `N`, both independently obtainable. No free parameters — **which is a strength: the postulate has nothing to tune and therefore nothing to overfit. PASS.**

Gate	Finding
FVP-6 FALSIFICATION	Two separable kill conditions. **(a) Functional form: fit `S = ln N . e^{−kΔ}` and test `k = 1`; a fitted k significantly different from 1 falsifies this postulate while leaving the framework standing. (b) Symmetry: fit separate exponents for `r > 1` and `r < 1`; if they differ, the symmetric form is wrong and the correction protocols of Chapter 48 must be reweighted accordingly.**
FVP-7 REPRODUCIBILITY	**UNTESTED.** Note that (a) and (b) above are cheap and could be run by any organisation with a decision log and a forecast record; this is the most tractable open test in Part IX.

FVP TABLE 47-E — The join `ln r = ln D − ln p_bar − F`

Gate	Finding
FVP-1 DIMENSIONS	Every term in nats after taking logs of dimensionless or commensurate quantities. **PASS.**
FVP-2 SCALE TYPE	Log-ratios are interval-scaled differences; the identity is an equation between such differences. **PASS.**
FVP-3 LIMITING CASES	`F −> 0`: `ln r −> ln(D/p_bar)`, maximum distortion for the given doctrine load. `F −> ln N`: `ln r −> ln(D/(p_bar N))`, minimum achievable. `D −> 0`: `ln r −> −inf`, Case 2. **PASS.**
FVP-4 FIXED POINTS	`F* = ln(D/p_bar)` gives `r = 1`, feasible iff `D/p_bar <= N`. **The infeasibility condition is the fixed point's domain of existence, and it is the chapter's hardest result. PASS.**
FVP-5 IDENTIFIABILITY	Given independent measurements of `D`, `p_bar` and `u`, the identity is over-determined — `ln r` can be computed two ways (directly from `D/P`, and via `ln D − ln p_bar − F`) and the two must agree. **This over-determination is a consistency check, and it is the most useful diagnostic in the chapter: a mismatch means one of the three measurements is wrong. PASS.**

Gate	Finding
FVP-6 FALSIFICATION	The identity itself is algebraic and cannot be falsified; what is falsifiable is the Aggregation Postulate it rests on. **Kill condition: measured P fails to track p_bar . N_eff across organisations with known u.**
FVP-7 REPRODUCIBILITY	**UNTESTED, and it inherits 47-C's status.**

47.11 What this chapter has and has not shown

It has shown that the moment equation, transposed to group scope, produces four limiting cases that each name a recognisable institutional condition; that its fixed point at `D = P` yields `M = E` and, as a derived consequence, exactly zero headroom; that Paper A's surplus enters the log-ratio as an exact subtractive term; and that the resulting `ln N` ceiling on correction reproduces, in the perceptual register, Paper E's `ln N` ceiling on sustainable growth.

It has not shown that any organisation behaves this way. Every FVP-7 row above reads UNTESTED, and that is the honest state of the work. What the chapter delivers is a specification precise enough that the test can be built: named quantities, declared units, a stated normalisation, an over-determination check, and four kill conditions.

And it has not shown that the numbers in §47.7 through §47.9 are measurements. Gate 2 is demoted, for the reasons set out at the head of §47.10: `D` and `P` are not yet on a common ratio scale, so the envelope of 0.53 to 1.90, the `u*` requirements and the `Δ_min` floors are illustrative constructions demonstrating the model's structure. The structure itself is untouched — the limits, the fixed point, the zero-headroom identity, the two-track ledger and the exact join `ln r = ln D − ln p_bar − F` all stand on ordinal information or on algebra. **A chapter that demotes its own headline envelope, states exactly which of its results survive, and prices the instrument that would restore the rest at 123 items per term is doing what the Formula Validation Protocol was written to make it do.**

The Specification

Requirements for measuring and managing a group's moment. Each names the term it moves.

G-1 . Declare the decision space and count `N`. Enumerate the distinct domains the organisation's intelligence must span, and justify each as genuinely distinct rather than relabelled. `[N]`
Test: an independent reader reproduces the list from the written definition.

G-2 . Adopt the commensurate-units convention explicitly. Write down, before measuring, what one effective evidence-channel per period means in this organisation. `[FVP–1]`

G-3 . Audit the channels for redundancy before counting them. Headcount and subscription count are not `P`. Compute `N_eff = 1 + (N–1)u` with `u` from error decorrelation on questions with later-observable answers. `[P, u]`

G-4 . Measure `p_bar` separately from `N_eff`. Per FVP 47-C/FVP-5, `P` alone does not identify them, and the two have entirely different fixes. `[P]`

G-5 . Enumerate the doctrine. List the decision-relevant propositions the standing model fixes in advance; weight each by the fraction of decisions it constrains; sum. That is `D`. `[D]` *Test: the list is written, dated, and each entry names the decisions it governs.*

G-6 . Run the over-determination check. Compute `ln r` twice — once as `ln D − ln P`, once as `ln D − ln p_bar − F` — and require agreement. A mismatch means one of the three measurements is wrong; find out which before proceeding. `[FVP 47–E/FVP–5]`

G-7 . Never report Δ without the sign of `ln r`. Per FVP 47-B/FVP-3, the two pathologies map to the same magnitude and require opposite corrections. `[reporting]`

G-8 . Never report `D` without an independent measurement of `P`. Per FVP 47-A/FVP-5, `D` is unidentifiable without the anchor. `[reporting]`

G-9 . Compute the infeasibility check first. If `D/p_bar > N`, stop: no arrangement of people calibrates this organisation. Record the residual floor `Δ_min = ln(D/p_bar) − ln N` and proceed to doctrine reduction (Ch 48). `[D, N]`

G-10 . Compute the operating envelope. From η, `C`, `N`, derive `S_min = C/η` and the band `min(r, 1/r) >= S_min/ln N`. That band is the organisation's licence to operate in this model. `[operating limits]` *Test: the two band edges are written down and dated.*

G-11 . Treat `E` as measurable and measure it. An organisation that never estimates the decision-relevant structure in its environment cannot distinguish Case 3 from Case 1 or 2. `[E]`

G-12 . Do not treat `r = 1` as an unqualified target. Per §47.6, exact calibration drives equilibrium headroom to zero. Decide deliberately whether the organisation is optimising banked order or preserving capacity for the new. `[h]`

G-13 . Justify the order proxy independently of the input. Paper E §6.3. The organisation's `Q` measure must not be a repackaging of its intelligence volume. `[Q]`

G-14 . Preregister before intervening. Encoder, fit window, freeze date, held-out window, four baselines including the shuffled-input control, and the kill conditions of FVP 47-A/FVP-6. `[falsifiability]`

G-15 . Re-measure `u` on a schedule. `u` decays (Ch 11, Ch 46). A calibrated organisation drifts into over-framing with no change in doctrine. Interval no longer than the measured half-life of `u`. `[u]`

G-16 . No protocol here may remove a person's ability to decline — to decline the measurement, to decline the channel role, or to decline the dissent it asks of them. CANON §5.4. `[ethics]`

The Failure Modes

FM-47.1 . THE VOID DENOMINATOR

Description. `P` computed as headcount or subscription count, without the redundancy audit. **Damages.** `P` overstated, therefore `r` understated; the organisation believes it is calibrated. **Signature.** Large, well-funded intelligence function; confident briefings; recurring surprises. **Measurement.** Error decorrelation across channels on questions with observable answers (G-3). **Fix.** Recompute `P` as `p_bar . N_eff`. Expect the number to fall by a large factor; at `u = 0.05` in an `N = 12` domain, `N_eff = 1.55` against a headcount that may be forty.

FM-47.2 . THE UNIDENTIFIED RATIO

Description. `D` reported without an independent `P`. Four distinct parameter sets produce the same observation (§47.10, FVP 47-A/FVP-5). **Damages.** measurement validity. **Signature.** Doctrine described in absolute terms — "our model is very strong" — with no channel count behind it. **Measurement.** Ask for `P`. If it is not measured, `D` is not measured. **Fix.** The anchor: count the audited channels, then `D = r . P`.

FM-47.3 . THE BARE MAGNITUDE

Description. `Δ` reported without the sign of `ln r`. **Damages.** diagnosis; the correction is applied in the wrong direction. **Signature.** A distortion score with no direction, followed by an intervention that makes things worse. Adding channels to an already under-framed organisation is the common instance (§47.8). **Measurement.** Read the report; is `ln r` signed? **Fix.** Always report the signed log-ratio. `Δ` is for magnitude comparison only.

FM-47.4 . THE CALIBRATED IRRELEVANCE

Description. `r = 1`, `Δ = 0`, and `E -> 0`. Perfect contact with a vanished environment. **Damages.** `M`, totally, while every internal metric reads well. **Signature.** Excellent internal process metrics; outputs that land nowhere; a constituency that has migrated. **Measurement.** Estimate `E` directly (G-11). Case 3 is distinguishable from Cases 1 and 2 only if `E` is measured. **Fix.** Outside the scope of this equation: the organisation must change what environment it is in. The formula's contribution is the diagnosis, not the remedy.

FM-47.5 . THE ZERO-HEADROOM SUCCESS

Description. The organisation calibrates, banks to its maximum, and finds `h = 0` exactly (§47.6). Every period is maintenance. **Damages.** `h`. **Signature.** A well-run institution that cannot start anything, with no identifiable culprit. **Measurement.** `h = η ln N − δQ − C` at the current `Q`, against `Q_inf = (η ln N − C)/δ`. **Fix.** Shed `Q` deliberately (Ch 23), or raise `η`, or cut `C`. Note that raising `N` is logarithmic and the worst buy. **Do not attempt to fix it by degrading calibration; that reduces S faster than it releases budget.**

FM-47.6 . THE PROXY THAT WAS THE INPUT

Description. The order proxy `Q` is a repackaging of intelligence volume — "we produced 400 assessments" counted as stored order. **Damages.** identifiability of `η`; the whole model becomes

untestable. **Signature.** Excellent training fit, held-out failure. η near 1 with negligible residual. **Measurement.** Paper E §6.3 check. **Fix.** An order proxy with an external referent: decisions whose outcomes are later observable, standing capability a third party can verify, infrastructure that survives staff turnover.

FM-47.7 . THE UNBADGED EQUATION

Description. `M = (D/P).E` presented as though it were a theorem of Papers A-E. **Damages.** the argument's warrant. **Signature.** "The mathematics shows this organisation is over-framed." **Measurement.** Read the sentence. Papers A-E prove nothing about `M`. **Fix.** Restate with the badge on: the moment equation is the author's postulate, its consequences are computable, and FVP-7 has not been run.

Worked Example — A Regional Authority

A constructed case carried end to end. **Every number is the example's own; none is drawn from a real organisation and none should be cited as empirical.** The arithmetic is reproducible from the stated parameters.

The system

A regional authority of about forty staff, responsible for infrastructure decisions across a defined territory. It reviews on a quarterly cycle. It has an intelligence function: it commissions surveys, maintains a monitoring programme, takes submissions, and reads the technical literature.

1. Encoding (Paper E §5.1's five required choices)

Choice	Setting
State encoder	Each intelligence channel's assessment vector across the 12 decision domains
Interaction observable	Cross-channel error decorrelation on questions whose answers become observable within 24 months
Order observable `Q`	Standing capability that survives staff turnover: verified asset records, executable procedures, precedent decisions upheld on review
Time scale	One quarter = one period
Cost observable `C`	Coordination, committee time, compliance, and the maintenance of the standing record

`Q` is measured by survival across turnover, which is external to intelligence volume, satisfying G-13 and Paper E §6.3.

2. Baseline measurement

```
N       = 12        ln N = 2.4849       twelve declared decision domains
u       = 0.14      measured by cross-channel error decorrelation
p_bar   = 1.00    one intake-unit per genuinely independent channel per quarter
D       = 5.20      doctrine load, in the same currency (G-5)
E       = 1.00      environment, normalised to one unit for legibility
eta     = 0.42      order-units per nat
delta   = 0.14      per quarter
C       = 0.55      order-units per quarter
```

Geometry (G-3):

```
N_eff = 1 + 11(0.14) = 2.5400
F     = ln(2.5400)   = 0.9322 nats
P     = p_bar . N_eff = 2.5400
```

Forty staff, an intelligence function with twelve declared domains, and an effective count of 2.54. That is the ratio-of-channels-to-headcount problem stated numerically: `2.54 / 40 = 0.0635`. The authority's audited intake is worth about six per cent of its headcount.

3. The ratio, computed two ways (G-6)

```
    direct:       ln r = ln D - ln P     = ln 5.20 - ln 2.5400 = 1.6487 -
0.9322 = 0.7165
      via the join: ln r = ln D - ln p_bar - F = 1.6487 - 0.0000 -
0.9322 = 0.7165     AGREE
```

The over-determination check passes, so the three measurements are mutually consistent.

```
    r     = e^{0.7165} = 2.0472
    Δ     = 0.7165 nats,  sign POSITIVE  ->  OVER-FRAMED  (G-7)
    M     = r . E = 2.0472
```

The authority constructs a situation about twice as large as its environment. It sees two units of structure where one exists. Behaviourally: it is confident, it is fast, and it is wrong in a direction its own model predicts.

4. The infeasibility check (G-9)

```
   D / p_bar = 5.20 / 1.00 = 5.20       N = 12       5.20 <= 12  ->  FEASIBLE
    u* = (5.20 - 1) / 11 = 0.3818         N_eff* = 5.20      F* = 1.6487
```

The authority *can* be calibrated by geometry alone. It needs `u = 0.3818`, which is `θ = 38.2` degrees of mean non-parallelism across its channels, against a measured 0.14 (`θ = 22.0` degrees). That is a large but not impossible move.

5. The operating envelope (G-10)

```
    S_min = C / eta = 0.55 / 0.42 = 1.3095 nats
    S_min / ln N    = 1.3095 / 2.4849 = 0.5270
    band: 0.5270 <= min(r, 1/r)   ->   0.527 <= r <= 1.898       Δ <= 0.6406
```

The authority's measured `r = 2.0472` is outside its own operating envelope. Not marginally — `Δ = 0.7165` against a limit of 0.6406. The model predicts its standing capability is eroding.

Confirming with the Contact Postulate and the storage law:

```
S     = ln N . min(r, 1/r) = 2.4849 x (1/2.0472) = 1.2138 nats
eta S = 0.42 x 1.2138 = 0.5098
Q_inf = (0.5098 − 0.55) / 0.14 = −0.2872
```

A negative equilibrium. Paper E Theorem 3.5 does not hide this failure mode and neither does the manual: as modelled, the authority's standing capability erodes to nothing with a time constant of `1/δ = 7.1` quarters, and the erosion will be attributed to turnover, budget, and the declining quality of submissions.

6. The two levers, priced

Lever A — raise `u` (recruit across lineages, rotate, commission from unaligned sources, stand up a genuinely independent channel):

u	N_eff	P	r	Δ	S	etaS	Q_inf
0.1400	2.540	2.540	2.0472	0.7165	1.2138	0.5098	−0.2872
0.2000	3.200	3.200	1.6250	0.4855	1.5292	0.6423	+0.6589
0.2400	3.640	3.640	1.4286	0.3567	1.7394	0.7306	+1.2897
0.3000	4.300	4.300	1.2093	0.1900	2.0548	0.8630	+2.2359
0.3818	5.200	5.200	1.0000	0.0000	2.4848	1.0436	+3.5259
0.4500	5.950	5.950	0.8739	0.1347	2.1717	0.9121	+2.5865

Two things to read. First, **the authority crosses back into its operating envelope at `u = 0.158`** — the value at which `N_eff = 5.20/1.898 = 2.7397` — well short of full calibration. It does not need to be perfect; it needs to be inside the band. Second, **the last row over-corrects.** At `u = 0.45` the authority passes calibration and begins under-framing, and `Q_inf` falls from 3.53 to 2.59. Diversity past the calibration point is a cost, exactly as §47.8 predicted.

Lever B — cut `D` (retire standing assumptions, sunset guidance, shed doctrine):

D	r	Δ	S	Q_inf
5.20	2.0472	0.7165	1.2138	−0.2872
4.50	1.7717	0.5719	1.4026	+0.2792
4.00	1.5748	0.4541	1.5779	+0.8052
3.50	1.3780	0.3206	1.8033	+1.4814
3.00	1.1811	0.1664	2.1039	+2.3831
2.54	1.0000	0.0000	2.4849	+3.5261

The authority re-enters its envelope at `D = 4.82`, since `2.5400 x 1.898 = 4.8209` — a reduction of about seven per cent in standing doctrine. **Both levers reach the same destination and the arithmetic is symmetric, because the Contact Postulate is symmetric in `ln r`.** Which to pull is not a question the formula answers; it is a question of cost, and Chapter 48 addresses it directly.

7. Headroom, and the trap at the end

```
Q      h = eta lnN − delta Q − C
----   ------------------------
0.0            +0.4937
1.0            +0.3537
2.0            +0.2137
3.0            +0.0737
4.0            −0.0663
5.0            −0.2063

h = 0 exactly at Q = 3.5261
```

Compare that to §47.6's derived result: `Q_inf` at perfect calibration is `(η ln N − C)/δ = 3.5261`, and `h = 0` at `Q = 3.5261`. **They are the same number, as the algebra requires.** If the authority calibrates fully and holds there, it will arrive at exactly zero headroom and will thereafter be able to maintain but never to build.

The recommendation the model supports is therefore not "calibrate fully". It is: get inside the envelope, stop short of the fixed point, and hold `Q` below `Q_inf` deliberately (G-12). Target `u ~ 0.24`, which gives `Δ = 0.36`, `Q_inf = 1.29`, and residual headroom at that store of `0.4937 − 0.14(1.29) = +0.31`. **Slightly distorted and solvent beats perfectly calibrated and saturated**, which is not a result anyone would have guessed from the equation's rhetoric and is a direct consequence of its arithmetic.

8. What would falsify this

Preregistered, per G-14 and FVP 47-A/FVP-6:

1. `r` measured quarterly for eight quarters; held-out prediction of decision quality (assessed at 24-month outcome observation) for quarters 9-16.
2. Four baselines: persistence; `E` alone; matched autoregressive on decision quality; and the shuffled-`r` control.
3. Preregistered kill conditions: **(a)** `E` alone predicts held-out decision quality as well as `(D/P).E`; **(b)** the shuffled-`r` control matches the fitted model; **(c)** the residual's sign is independent of the sign of `ln r`; **(d)** for the Contact Postulate specifically, a fitted exponent `k` in `S = ln N . e^{−kΔ}` significantly different from 1, or significantly different between the `r > 1` and `r < 1` branches.

Kill condition (d) is separable from the rest and is worth stating on its own, because it is the cheapest test in Part IX and it can falsify one postulate while leaving the framework intact. **A framework that cannot fail cannot be trusted; this one names four ways to fail and one of them costs a decision log and a forecast record.**

The moment equation survives the change of scope. Its four limits name four institutions anyone can recognise, its fixed point yields undistorted contact and — as a derived surprise — exactly zero headroom, and Paper A's surplus enters its log-ratio as an exact one-for-one term, which is the join between Volume Two's two equations. What it does not yet have is a single measured

organisation. Chapter 48 takes the two sides of `r != 1` and specifies, for each, how to find out which one you are in and what to do about it.

Chapter 48 . Doctrine-to-Data

Chapter 47 established the group moment and its calibration ratio. This chapter is about the two ways an organisation gets that ratio wrong, and it is the most directly usable chapter in Part IX, because both failures are common, both are measurable with instruments an organisation can build from records it already keeps, and both are correctable.

They are not, however, correctable by the same means. **The two pathologies map to the same distortion magnitude and require opposite interventions, which is why an unsigned diagnosis is worse than no diagnosis at all.** An organisation told it is "distorted" and not told in which direction will reach for the nearest available remedy — usually more collection — and half the time that remedy is the disease.

The chapter closes a loop opened much earlier. Chapter 11 described pseudo-diversity: an institution whose members are epistemically parallel, whose agreement is read as convergence, and whose effective count is one. Chapter 11 priced that in surplus. This chapter shows what it does to perception, and the result is stronger than the analogy suggests:

Pseudo-diversity is not merely correlated with over-framing. It is a mechanism that manufactures over-framing, from constant doctrine, with no one adding a single new belief.

The Law

48.1 The two directions, stated

From Chapter 47, carrying its badges forward:

```
   r = D / P        Δ = | ln r |        M = r . E        S = ln N . min(r, 1/r)

    r > 1   OVER-FRAMED     M > E    doctrine outruns intelligence
    r = 1   CALIBRATED      M = E    undistorted contact
    r < 1   UNDER-FRAMED    M < E    intelligence outruns doctrine
```

Status. `M = (D/P).E`, the Commensurate-Units Convention, the Aggregation Postulate `P = p_bar.N_eff`, and the Contact Postulate `S = ln N . e^{-Δ}` are all **[POSTULATES]**, stated and badged in Chapter 47 §47.1-47.4. `N_eff = 1 + (N-1)u` and `0 <= F <= ln N` are **[PROVEN]** (Paper A Thm 2.1; Paper D Cor 2.2). Nothing in this chapter upgrades a postulate.

48.2 The asymmetry that governs everything here

Δ is a magnitude and discards the sign. That is fine for comparing severity and fatal for choosing a treatment, because the two conditions sit on opposite sides of a fixed point and every lever moves `ln r` in one direction only.

```
     LEVER                                effect on ln r      helps OVER-
framed    helps UNDER-framed
    -------------------------   ----------------   ------------------   --
---------------
    raise u (more independence)   decreases         YES                NO —
makes it worse
```

```
    raise N (broader domain)      decreases          YES (weakly, log)    NO —
makes it worse
    raise p_bar (deeper intake)   decreases          YES                  NO —
makes it worse
    cut D (shed doctrine)         decreases          YES                  NO —
makes it worse
    raise D (build a frame)       increases          NO — makes it worse  YES
    triage P (cut channels)       increases          NO — makes it worse  YES
```

Every standard organisational virtue in the left column — more diversity, more breadth, more collection, less dogma — is a poison to an under-framed institution. Chapter 47 §47.8 computed this: an organisation carrying `D/p_bar = 6.0` at `N = 12` reaches `r = 0.5` at `u = 1.0`, and its bankable input falls from the ceiling of 2.4849 nats back to 1.2425. Diversity past the calibration point is a cost, and the model says so without hedging.

This is the chapter's governing rule and it is stated as a requirement, not a suggestion:

> **NEVER REPORT A DISTORTION WITHOUT ITS SIGN.** Report `ln r`, signed. Δ exists for comparing severity across units, never for prescribing.

48.3 The pseudo-diversity mechanism, stated at law level

Chapter 11's finding: `u` is not a property an institution acquires but a quantity that decays and must be maintained. Chapter 46 §46.15 showed the same decay in the two-person case and gave it a half-life. Put that decay into Chapter 47's join and something specific follows.

```
ln r_t = ln D − ln p_bar − F_t          F_t = ln( 1 + (N−1) u_t )

with   u_t = u_0 (1 − λ)^t     (Chapter 11's decay, Chapter 46's λ)
```

Hold `D` fixed. Hold `p_bar` fixed. Hold `N` fixed. Hold headcount, budget and collection volume fixed. **`ln r` still rises, monotonically, forever, because `F_t` falls.**

> **THE MANUFACTURING RESULT.** An organisation that adds no doctrine, cuts no collection, and changes no policy will drift from calibration into over-framing at a rate set entirely by the decay of `u`. The crossing occurs at
>
> `u∗ = ( D/p_bar − 1 ) / ( N − 1 )` `t∗ = ln( u∗ / u_0 ) / ln( 1 − λ )`
>
> and after `t∗` the organisation is over-framed, with an unchanged model, by arithmetic alone.

Worked, at `N = 12`:

D/p_bar	u_0	λ	u∗	t∗ (periods)
4.0	0.55	0.08	0.2727	8.41
4.0	0.55	0.04	0.2727	17.18
3.0	0.30	0.06	0.1818	8.09

Roughly two years, at quarterly periods and an eight-per-cent-per-quarter convergence rate, from calibrated to over-framed with nobody doing anything. Halving the decay rate roughly doubles the time, as the log form requires; it does not prevent the crossing.

This is why the connection to Chapter 11 is a mechanism and not an analogy. Pseudo-diversity's signature — `u -> 0`, `N_eff -> 1`, surplus `ln(1) = 0` — enters `r = D/(p_bar.N_eff)` as a *collapsing denominator*. **The institution's doctrine did not grow. Its denominator vanished.**

48.4 The honesty clauses

1. Every instrument in this chapter estimates `r` **through a proxy**, and each proxy carries its own badged postulate linking the observable to the ratio. None of them measures `D` or `P` directly.
2. Chapter 47's FVP 47-A/FVP-5 finding stands: **only `r` is identifiable from `M` and `E`**; `D` requires an independent anchor on `P`.
3. The Contact Postulate's symmetry — over- and under-framing attenuating equally at equal `Δ` — is **untested** and is the first thing a real study should attack (Ch 47, FVP 47-D/FVP-6).
4. `Δ` values in this chapter are model quantities. **No empirical distribution of institutional `Δ` exists**, because nobody has measured one. Where a threshold is recommended below, it is derived from the organisation's own `S_min = C/η`, not from a literature.
5. FVP-7 is unrun for every formula here.

The Engineering

48.5 The over-framed institution

Clinical description. `r > 1`. The organisation's standing model supplies more constraint on its decisions than its audited intelligence does. It therefore constructs a situation larger and more determinate than its environment: `M = rE > E`. It sees what its model predicts.

The condition is not stupidity and it is not usually arrogance. It is the ordinary end-state of a successful institution: doctrine accumulates because doctrine is what an institution is *for* (Chapter 23 — `Q` at civilisational scale), while `u` decays because convergence is what collaboration produces (Chapter 46 §46.15).

The behavioural signature, in the order it usually appears:

1. Assessments become faster and more confident.
2. New evidence arrives and is *absorbed* — explained within the model — rather than escalated.
3. Conclusions stop moving. The revision rate falls.
4. Surprises arrive from outside the model's categories, are treated as one-offs, and are absorbed.
5. Internal metrics improve while external outcomes do not.

Point 5 is the diagnostic crux. **An over-framed organisation's self-assessment improves as its condition worsens**, because self-assessment is conducted through `M`, and `M` is exactly what is inflated. This is why the instruments below are all built on records with an *external* referent.

48.6 Diagnostic instruments for over-framing

Four instruments. Each states its postulate, its data requirement, its scale type, and what it can and cannot conclude.

I-1 . THE REVISION RATE

> **THE REVISION POSTULATE [POSTULATE].** A calibrated organisation revises standing conclusions at some baseline rate σ∗ determined by the volatility of its environment. The observed revision rate σ — the fraction of intelligence products that alter a standing conclusion — estimates the ratio by r̂ = σ∗ / σ.

Data: a decision or assessment log with, per product, a recorded before-state and after-state of the standing conclusion. Scale: σ is a proportion, ratio-scaled with a true zero.

σ	r̂	ln r̂	reading
0.40	0.500	−0.6931	under-framed
0.30	0.667	−0.4055	under-framed
0.20	1.000	+0.0000	calibrated (σ∗ = 0.20 here)
0.15	1.333	+0.2877	over-framed
0.10	2.000	+0.6931	over-framed
0.05	4.000	+1.3863	over-framed, severe
0.02	10.000	+2.3026	over-framed, severe

Limits, stated plainly. σ∗ is not known a priori and must be estimated from the environment's own volatility — the rate at which the decision-relevant facts actually change — which requires an external referent (audit outcomes, realised events, later-observable answers). **Without a defensible σ∗ this instrument produces an ordering, not a magnitude**, and must be reported as such. It is still useful in that demoted form: σ falling over time, with a stable environment, is strong evidence of drift toward over-framing.

I-2 . THE ANOMALY DISPOSAL AUDIT

Take every observation over a period that was inconsistent with the standing model. Classify each into three bins: **escalated** (it changed something), **absorbed** (it was explained within the model), **discarded** (it was set aside as noise, error, or an outlier).

A = (absorbed + discarded) / total anomalies escalation rate = 1 − A

A	escalation	odds absorb:escalate	ln odds
0.30	0.70	0.43 : 1	−0.8473
0.50	0.50	1.00 : 1	+0.0000
0.70	0.30	2.33 : 1	+0.8473
0.80	0.20	4.00 : 1	+1.3863
0.90	0.10	9.00 : 1	+2.1972
0.95	0.05	19.00 : 1	+2.9444
0.98	0.02	49.00 : 1	+3.8918

The log-odds is the right register — it is a log-ratio, commensurate with ln r and with F. **This instrument is the most robust of the four**, because it requires no baseline: the raw classification counts are meaningful on their own, and the audit is conducted on records the organisation already generates whenever it writes down why something was not acted on.

Its weakness is that anomaly identification is itself performed through the model. **An organisation severely over-framed does not record anomalies as anomalies; it does not see them.** The audit therefore *understates* severe cases, which makes a high A conclusive and a low A inconclusive. Say so in the report.

I-3 . THE CALIBRATION CURVE

Plot stated confidence against realised frequency across a forecast record. Over-framed institutions are overconfident: the 90-per-cent band contains the outcome less than 90 per cent of the time. Under-framed institutions are the reverse, or — much more commonly — **make no quantified forecasts at all**, which is itself the diagnosis (see I-8).

Data requirement: a forecast register with resolution dates and recorded outcomes. Scale: proportions on both axes, ratio-scaled.

This instrument does not estimate r numerically without a further postulate mapping overconfidence to the ratio, and this manual declines to invent one. **It is reported as corroborating evidence with a direction, not as a magnitude.** Its value is that it is entirely external: nothing about it passes through M.

I-4 . THE CHANNEL AUDIT — the structural instrument

The other three are behavioural. This one measures the mechanism directly, and it is the instrument the manufacturing result of §48.3 makes indispensable.

```
   Step 1.  Enumerate the intelligence channels and the N declared domains.     (Ch 47, G-
1)
    Step 2.  For questions whose answers later become observable, record each
             channel's assessment BEFORE any cross-channel discussion.
   Step 3.  Compute pairwise error decorrelation; take the mean as u.         (Ch 9)
    Step 4.  N_eff = 1 + (N-1)u.    F = ln N_eff.    P = p_bar . N_eff.
   Step 5.  ln r = ln D - ln p_bar - F, with D from the doctrine enumeration.    (Ch 47, G-
5)
   Step 6.  Cross-check against ln D - ln P.  They must agree.                (Ch 47, G-
6)
```

Worked at N = 12, showing what headcount hides:

u	N_eff	F	P/headcount at 40 staff
0.00	1.000	0.0000	0.0250
0.02	1.220	0.1989	0.0305
0.05	1.550	0.4383	0.0387
0.10	2.100	0.7419	0.0525
0.15	2.650	0.9746	0.0663
0.25	3.750	1.3218	0.0938
0.40	5.400	1.6864	0.1350
0.60	7.600	2.0281	0.1900

A forty-person intelligence function at `u = 0.05` has an audited intake of 1.55 channels — under four per cent of its headcount. That is the number that goes in the denominator, and it is the number nobody computes.

The channel audit is the only instrument of the four that tells you *why*. The other three tell you that `r` is high; this one tells you whether `D` grew or `P` collapsed, and those have different fixes.

48.7 The under-framed institution

Clinical description. `r < 1`. Audited intelligence exceeds the constraint the standing model supplies. The organisation constructs a situation smaller and less determinate than its environment: `M = rE < E`. It cannot make its inputs mean anything.

This is the less-discussed pathology and it is at least as common. It is the natural state of a new organisation, a merged one, a rapidly expanded one, and any organisation that responded to a past over-framing episode by collecting more.

The behavioural signature:

1. Intake grows. Dashboards multiply. Reporting proliferates.
2. Decision latency grows. Items are deferred pending further information.
3. The standing request is always for more data.
4. Nothing is retired: no conclusion is ever firm enough to close a line of enquiry.
5. Staff describe the work as overwhelming and describe the organisation as lacking direction. Both reports are accurate and they are the same report.

Note that under-framing is **not** ignorance. It is the opposite: a surfeit of intake with insufficient frame to convert intake into a situation. CANON_INTERNAL §4's phrase for the personal case applies unchanged at group scope — **sensation is not experience** — and Chapter 47's Case 2 gives the limit exactly: `D -> 0` implies `M -> 0`, not `M` small but `M` zero. A perfectly open-minded institution does not decide badly. It does not decide.

48.8 Diagnostic instruments for under-framing

I-5 . THE INTAKE-TO-DECISION RATIO

> **THE CONVERSION POSTULATE [POSTULATE].** Let κ be intake-units consumed per decision closed. A calibrated organisation runs at some κ∗ set by its decision complexity. Then r̂ = κ∗ / κ.

```
κ / κ∗     r̂        ln r̂
-------   ------    --------
  0.25    4.0000    +1.3863     over-framed: decides on almost nothing
  0.50    2.0000    +0.6931
  1.00    1.0000    +0.0000     calibrated
  2.00    0.5000    -0.6931
  4.00    0.2500    -1.3863     under-framed
  8.00    0.1250    -2.0794     under-framed, severe
```

Data: an intake tally and a decision log, both of which most organisations keep. Scale: both are counts, so κ is ratio-scaled. Same limitation as I-1: $\kappa*$ requires justification, and without one the instrument produces an ordering rather than a magnitude. **Trend in κ under stable decision complexity is the usable form.**

I-6 . THE UNREAD RATIO AND THE BACKLOG SLOPE

Two counts: intelligence products produced, and intelligence products demonstrably consulted in a decision within their useful life. The ratio is a direct, external, ratio-scaled measure of intake the organisation cannot frame.

A rising backlog with a stable production rate is Little's Law read backwards: $L = \lambda W$, so a growing queue L at constant arrival rate λ means growing residence time W. **Under-framing shows up in the queue before it shows up in anyone's account of the problem**, which makes this the earliest available indicator.

I-7 . THE INSTRUMENT PROLIFERATION COUNT

Count distinct standing reports, dashboards and metrics; count how many are cited in decisions. Proliferation with flat citation is diagnostic: the organisation is answering an inability to frame by adding intake, which moves r in exactly the wrong direction.

This instrument has a specific value: **it detects the characteristic self-inflicted worsening of under-framing.** An under-framed organisation's instinct is to collect more, and collecting more lowers r further. The count catches the loop.

I-8 . THE FORECAST-ABSENCE TEST

Ask for the organisation's quantified forecasts with resolution dates. An over-framed organisation has them and they are overconfident (I-3). An under-framed organisation **does not have them at all** — not because it is careless, but because a forecast is an act of doctrine, and it has insufficient doctrine to make one.

Binary, external, and cheap. Its absence is not proof, but combined with I-5 and I-6 it is close to conclusive on direction.

48.9 Reading the instruments together

No single instrument is sufficient. The direction table:

		OVER-FRAMED	UNDER-FRAMED
I-1	revision rate σ	LOW	HIGH
I-2	anomaly absorption A	HIGH	LOW (or unclassifiable)
I-3	calibration curve	overconfident	underconfident / absent
I-4	channel audit u, N_eff	LOW u, LOW N_eff	HIGH N_eff relative to D
I-5	intake per decision κ	LOW	HIGH
I-6	unread ratio, backlog	LOW	HIGH, rising
I-7	instrument proliferation	LOW	HIGH with flat citation

I-8	quantified forecasts	present, wrong	absent

Rules for reading:

- **Direction requires at least two instruments agreeing, one behavioural and one structural.** I-4 should be one of them wherever the data allow, because it is the only one that identifies the cause.
- **A mixed reading is not noise; it is usually a mixed organisation.** See §48.11.
- **I-2 is conclusive when high and inconclusive when low** (§48.6). Never read a low A as evidence of calibration.
- **Where σ∗ and κ∗ cannot be defended, report I-1 and I-5 as orderings, not magnitudes**, per FVP-2 and Chapter 53. The demotion must appear in print.

48.10 The manufacturing result, computed end to end

Take an organisation calibrated on the under-framed side and let nothing happen except u decay. N = 12, D = 4.0, p_bar = 1.0, u_0 = 0.55, λ = 0.08 per period, and for the storage side η = 0.42, δ = 0.14, C = 0.55.

t	u	N_eff	F	P	r	ln r	Δ	S	Q_inf
0	0.5500	7.050	1.9530	7.050	0.567	−0.5667	0.5667	1.4099	+0.3011
3	0.4283	5.711	1.7424	5.711	0.700	−0.3561	0.3561	1.7404	+1.2927
6	0.3335	4.668	1.5408	4.668	0.857	−0.1545	0.1545	2.1291	+2.4588
9	0.2597	3.857	1.3498	3.857	1.037	+0.0365	0.0365	2.3958	+3.2589
12	0.2022	3.224	1.1707	3.224	1.241	+0.2156	0.2156	2.0031	+2.0806
15	0.1575	2.732	1.0051	2.732	1.464	+0.3812	0.3812	1.6973	+1.1632
18	0.1226	2.349	0.8539	2.349	1.703	+0.5324	0.5324	1.4591	+0.4488
21	0.0955	2.050	0.7180	2.050	1.951	+0.6683	0.6683	1.2737	−0.1075
24	0.0743	1.818	0.5976	1.818	2.200	+0.7887	0.7887	1.1293	−0.5407

The crossing is at u∗ = (4.0 − 1)/11 = 0.2727, reached at t∗ = ln(0.2727/0.55)/ln(0.92) = 8.41 periods.

Read the trajectory as a story, because it is the commonest institutional history there is.

- **Periods 0-6.** The organisation is under-framed and improving. Its distortion falls, its bankable input rises, its equilibrium store rises. Everything anyone measures is getting better.
- **Period 9.** Calibration. Δ = 0.037, S = 2.3958 against a ceiling of 2.4849. Peak performance, and it is genuinely peak: Q_inf = 3.26.
- **Periods 12-18.** The same process continues and now runs the other way. u keeps falling, F keeps falling, ln r keeps rising, and every quantity that improved now degrades — **driven by the identical underlying trend, with no change of behaviour, no change of doctrine, and no decision by anyone.**
- **Period 21.** Q_inf goes negative. The institution's standing capability now erodes.

The mechanism that produced the golden period is the mechanism that produced the decline. They are the same mechanism at different points on one monotone curve.

This is why organisations are so reliably unable to diagnose themselves at period 18. The behaviour that was rewarded at period 6 — building shared understanding, converging on a common view, reducing internal friction — is the behaviour that is killing them at period 18, and it never changed. **Pseudo-diversity manufactures over-framing, and it does so most effectively in organisations that are proud of their cohesion.**

Chapter 18's warning arrives here with an equation attached: cohesion is often the enemy of the collective-intelligence factor because cohesion drives u toward zero. In this chapter's terms, cohesion drives the denominator of `r` toward `p_bar`.

48.11 Distortion does not average

Organisations are not uniform. A headquarters and a field operation can sit on opposite sides of the fixed point, and the aggregate figure will report neither.

Take two units, `N = 12`, `p_bar = 1.0` throughout:

unit	u	N_eff = P	D	r	ln r	Δ	S
HQ	0.11	2.210	6.80	3.0769	+1.1239	1.1239	0.8076
Field	0.62	7.820	1.15	0.1471	−1.9169	1.9169	0.3654

AGGREGATE (D and P summed): 7.95 / 10.030 = 0.7926 ln r = −0.2324

The aggregate reads `Δ = 0.2324` — nearly calibrated — while both halves are severely distorted, at 1.12 and 1.92. The aggregate is not an average of the units' distortions and it is not a summary of them; it is a different quantity, and here it is actively misleading.

The arithmetic of why: `r_agg = (D_1+D_2)/(P_1+P_2)` is a weighted mediant, so it always lies between `r_1` and `r_2`. When both units are distorted in the **same** direction, the aggregate lies between them and understates the worse one — tolerable. When they are distorted in **opposite** directions, the aggregate lies between two values straddling 1, and can sit arbitrarily close to 1 while both units are arbitrarily far from it. **Cancellation in the books, no cancellation in the world.**

The consequence for banking is severe. Using the Contact Postulate:

```
S computed from the aggregate ratio:            2.4849 x 0.7926 = 1.9696
S computed per unit and intake-weighted:          0.4629
OVERSTATEMENT FACTOR:                             4.26 x
```

Computing an organisation's bankable input from its aggregate ratio overstates it by a factor of 4.26 in this case. The aggregate is not conservative and not approximately right; it is wrong in the dangerous direction.

> **RULE. Compute `r` per decision-making unit. Aggregate the units' `S`, weighted by intake, never the units' `D` and `P`.**

And one further consequence, which is where §48.12's correction protocols start: a mixed organisation has an obvious-looking remedy — move doctrine from the over-framed unit to the under-framed one — and it does not fully work. Doctrine total is 7.95 and intake total is 10.030. **Even perfect redistribution leaves the organisation doctrine-poor overall by 2.080 units, or 26.2 per cent of its standing doctrine.** Redistribution alone can equalise the distortion; it cannot remove it.

48.12 Correction protocols — OVER-FRAMED

The correction runs in one direction only: `ln r` must come down. There are exactly four levers and they have very different prices.

Lever O-1 . Raise `u`. Recruit across training lineages, rotate, stand up genuinely independent channels, commission from unaligned sources, install adversarial review, import a practitioner from an unrelated trade (Chapter 9's full inventory).

Price and ceiling. `ln r` falls by exactly the gain in `F`, one nat for one nat, and the total available is `ln N` (Chapter 47, Consequence 1). Concavity means the early gains are the large ones: at `N = 12`, moving `u` from 0.14 to 0.30 buys 0.5264 nats; moving from 0.60 to 0.90 buys 0.3607. **Raise `u` first, and expect the second independent channel to be worth more than the twentieth.**

The trap. `u` gained is `u` that will decay again at rate `λ`. This lever must be operated as a maintenance schedule, not a project (Chapter 46 §46.17, Chapter 47, G-15). A one-off diversity intervention buys a decaying quantity and is the single commonest wasted correction.

Lever O-2 . Raise `p_bar`. Deepen the intake of channels you already have: longer collection, better instrumentation, more careful reading.

Price. Directly proportional and unbounded in principle — `ln r` falls by `ln` of the multiple. In practice it is bounded by budget and by the fact that deepening a channel does not make it independent. **`p_bar` and `u` are not substitutes: a hundred times more of one wire service is still one channel.**

Lever O-3 . Raise `N`. Broaden the declared domain.

Price. Logarithmic and expensive (Chapter 8). It raises the ceiling on what `u` can buy but buys nothing by itself. **Worst of the four, and it is the one most often reached for, because "expanding our scope" is easier to authorise than "we are wrong".**

Lever O-4 . Cut `D`. Retire standing assumptions. Sunset guidance. Shed doctrine.

Price. Unbounded — `D` can go to zero — but it costs `Q`, and shedding `Q` is Chapter 23's move, which is legitimate and underused and hard. **This is the only lever that works in the infeasible region `D/p_bar > N`,** where no arrangement of people can calibrate the organisation.

The comparison, at Chapter 47's worked parameters (`N = 12`, `D = 5.2`, `p_bar = 1.0`, `u = 0.14`):

```
target r     via u                      via D
--------     ----------------------     ----------------------
1.898        u = 0.1582  (x 1.13)       D = 4.821  (−7.3 %)
1.500        u = 0.2242  (x 1.60)       D = 3.810  (−26.7 %)
1.200        u = 0.3030  (x 2.16)       D = 3.048  (−41.4 %)
```

```
1.000        u = 0.3818  (x 2.73)       D = 2.540  (−51.2 %)
```

Getting merely inside the operating envelope costs a 13 per cent increase in u or a 7 per cent cut in doctrine. Full calibration costs nearly a tripling of u or halving the standing model. The lesson is the same one Chapter 26 gives for interventions generally: stabilise before optimising, and take the cheap fraction of the distance first.

The sequence, as a procedure.

```
O-STEP 1.  Confirm direction with two instruments, one of them I-4.
O-STEP 2.  Compute the infeasibility check: is D/p_bar > N?
           If YES -> only O-4 works. Go to step 5.
O-STEP 3.  Compute the operating envelope from S_min = C/eta.
           Target the envelope edge, not r = 1 (Ch 47, G-12).
O-STEP 4.  Price O-1 and O-2 against O-4 at that target. Take the cheapest.
O-STEP 5.  If O-4: enumerate the doctrine (Ch 47, G-5), rank by
      (decisions constrained) / (evidence supporting), retire from the top.
O-STEP 6.  Install the maintenance schedule for u BEFORE declaring the
           correction complete. Without it, t* restarts.
O-STEP 7.  Preregister the held-out test (Ch 47, G-14) before intervening.
```

Step 5's ranking deserves a note. The doctrine to retire first is not the doctrine most likely to be wrong; it is the doctrine that constrains the most decisions on the least evidence. **A weakly supported proposition that governs one decision a year is a curiosity. A weakly supported proposition that governs every decision is the whole problem.**

48.13 Correction protocols — UNDER-FRAMED

The correction runs the other way and it is **not** the mirror image, for a reason that matters: building a frame is a slower, more skilled, and more contestable act than adding a channel.

Lever U-1 . Raise `D` — build the frame. Adopt an explicit model. Write the standing assumptions down. Commit to a hypothesis before collecting. Publish quantified forecasts with resolution dates. Close lines of enquiry.

Price. This is the only lever that addresses the cause, and it is expensive in the one currency institutions guard most closely: **someone must be willing to be wrong in writing.** Under-framing persists in organisations not because frames are hard to obtain but because an unstated frame cannot be falsified and a stated one can.

At $N = 12$, $u = 0.55$ so $P = 7.05$, $\eta = 0.42$, $\delta = 0.14$, $C = 0.55$:

D	r	Δ	S	Q_inf
0.50	0.0709	2.6462	0.1762	−3.3999
1.00	0.1418	1.9530	0.3525	−2.8712
2.00	0.2837	1.2599	0.7049	−1.8138
3.00	0.4255	0.8544	1.0574	−0.7564
4.00	0.5674	0.5667	1.4099	+0.3011
5.00	0.7092	0.3436	1.7623	+1.3585
6.00	0.8511	0.1613	2.1148	+2.4159
7.05	1.0000	0.0000	2.4849	+3.5261

This organisation's equilibrium store crosses zero between `D = 3` and `D = 4`. Below that it cannot hold anything, no matter how good its intelligence is. The remedy is not better collection. It has excellent collection. **It needs a model, and until it has one its collection is unbankable.**

Lever U-2 . Triage `P` — cut channels deliberately.

The uncomfortable lever, and it is real. Reducing the number of independent channels raises `r` toward 1. Worked at `D = 1.6`, `N = 12`:

u	P	r	Δ	S
0.5500	7.050	0.2270	1.4830	0.5640
0.4500	5.950	0.2689	1.3134	0.6682
0.3500	4.850	0.3299	1.1090	0.8198
0.2500	3.750	0.4267	0.8518	1.0602
0.1500	2.650	0.6038	0.5046	1.5003
0.0545	1.599	1.0003	0.0003	2.4841

Reaching calibration by triage alone would require destroying almost all of this organisation's independence — `u` from 0.55 to 0.055 — and it would work, in the sense that `S` returns to the ceiling. **The manual states the lever and then states why it is nearly always the wrong one:**

1. It is irreversible on the timescale that matters. `u` destroyed is `u` that must be rebuilt from recruitment, and Chapter 46 §46.15 shows that rebuilding difference has no internal source.
2. It caps the organisation's future. `N_eff` is the ceiling on every later correction; an organisation that triages its way to calibration at `N_eff = 1.6` has `D/p_bar <= 12` still, but has thrown away the geometry it would need if `D` ever grows.
3. It is the same move as pseudo-diversity, chosen deliberately. Everything §48.10 says about the trajectory then applies, starting immediately.

Use triage only for genuinely redundant channels — channels whose measured pairwise `u` against existing ones is near zero. Those cost `p_bar`-weighted budget and contribute nothing to `N_eff`, so cutting them raises `r` almost not at all and frees resources for U-1. That is the honest, narrow, defensible version of this lever.

The sequence, as a procedure.

```
      U-STEP 1.  Confirm direction with two instruments; I-8 and I-
6 are the cheapest pair.
   U-STEP 2.  Compute the operating envelope from S_min = C/eta and locate the
              D at which Q_inf crosses zero. That is the minimum viable frame.
    U-STEP 3.  Build the frame: write the standing assumptions, commit to
               hypotheses, publish forecasts with resolution dates.
    U-STEP 4.  Close lines of enquiry explicitly. An unclosed enquiry is
               continuing intake against a frame that does not yet exist.
    U-STEP 5.  Triage ONLY channels with measured near-zero pairwise u.
    U-STEP 6.  Re-measure. Under-framed organisations correct FAST once the
               frame lands, because D is the numerator and moves directly.
    U-STEP 7.  Set the over-framing alarm now: record t* from the current u,
               lambda and D, and diarise the re-measurement before it.
```

Step 7 is the step that closes the loop of this chapter. An organisation that corrects under-

framing by building doctrine has raised D — and D does not decay, while u does. It has therefore shortened its own t∗ and set itself on the trajectory of §48.10. The correction for one pathology is the seed of the other, and the only defence is to have computed the crossing date in advance.

48.14 FVP tables

FVP TABLE 48-A — Revision-rate estimator r̂ = σ∗/σ (I-1)

Gate	Finding
FVP-1 DIMENSIONS	σ, σ∗ are proportions (revisions per product), dimensionless; r̂ dimensionless. **PASS.**
FVP-2 SCALE TYPE	Proportions are ratio-scaled with a true zero (no revisions). **PASS for σ. The gate binds on σ∗:** if the baseline cannot be estimated from an external referent, r̂ has no magnitude and **the instrument must be demoted in print to an ordering.**
FVP-3 LIMITING CASES	σ –> 0: r̂ –> inf, an organisation that never revises — real, nameable, and the severe end of over-framing. σ –> 1: every product overturns a conclusion, r̂ –> σ∗, the maximally unstable frame. **PASS.**
FVP-4 FIXED POINTS	σ = σ∗ gives r̂ = 1. Not an attractor; nothing drives σ toward σ∗. **PASS, with instability noted.**
FVP-5 IDENTIFIABILITY	σ is recoverable from a decision log with before/after states. σ∗ is **not recoverable from the organisation's own records** and requires environmental volatility measured externally. **CONDITIONAL PASS.**
FVP-6 FALSIFICATION	Preregister σ over a fit window; predict held-out surprise rate — the frequency of outcomes outside the organisation's stated bands. **Kill condition: σ fails to predict held-out surprise rate beyond the shuffled-input control.**
FVP-7 REPRODUCIBILITY	**UNTESTED.** Cross-institution recovery of σ∗ for comparable environments is the specific open question.

FVP TABLE 48-B — Anomaly absorption A, reported as ln(A/(1–A)) (I-2)

Gate	Finding
FVP-1 DIMENSIONS	`A` a proportion; the log-odds in nats, commensurate with `ln r` and `F`. **PASS.**
FVP-2 SCALE TYPE	Counts of anomalies by disposal class: ratio-scaled with a true zero. **PASS, and this is the strongest FVP-2 position of any instrument in the chapter, because it needs no baseline.**
FVP-3 LIMITING CASES	`A -> 1`: nothing is ever escalated, log-odds `-> +inf`; the sealed model. `A -> 0`: everything is escalated, log-odds `-> -inf`; a frame so weak that no observation can be placed. Both real. **PASS.**
FVP-4 FIXED POINTS	`A = 0.5`, log-odds 0. Note this is **not** claimed to correspond to `r = 1`; the mapping from `A` to `r` is not established and this manual does not invent one. **PASS as a fixed point of the instrument, NOT of the ratio.**
FVP-5 IDENTIFIABILITY	`A` recoverable from disposal records. **The known bias, printed: anomaly identification is performed through the model, so severe over-framing suppresses the denominator. High `A` is conclusive; low `A` is not. CONDITIONAL PASS with declared bias.**
FVP-6 FALSIFICATION	Preregister `A`; predict the rate of later-observable events the organisation had classified as noise. **Kill condition: no relation between `A` and the missed-event rate beyond the shuffled control.**
FVP-7 REPRODUCIBILITY	**UNTESTED.** Inter-rater agreement on the three disposal classes is a prerequisite and has not been established.

FVP TABLE 48-C — Intake-to-decision estimator $\hat{r} = \kappa^*/\kappa$ (I-5)

Gate	Finding
FVP-1 DIMENSIONS	κ in intake-units per decision; κ^* the same; $\hat{r}$ dimensionless. **PASS, provided intake-units are the Chapter 47 convention's units.**
FVP-2 SCALE TYPE	Both are ratios of counts, ratio-scaled. **PASS for κ; binds on κ^* exactly as 48-A binds on σ^*. Demote to an ordering where κ^* is undefended.**

Gate	Finding
FVP-3 LIMITING CASES	κ –> 0: decisions taken on no intake, r̂ –> inf, over-framing. κ –> inf: unbounded intake per decision, r̂ –> 0, the paralysed institution of Chapter 47's Case 2. **PASS.**
FVP-4 FIXED POINTS	κ = κ∗. **PASS.**
FVP-5 IDENTIFIABILITY	κ from intake tally and decision log, both routinely kept. κ∗ requires an external standard of decision complexity. **CONDITIONAL PASS.**
FVP-6 FALSIFICATION	Preregister κ; predict held-out decision latency and reversal rate. **Kill condition: no relation beyond the shuffled control.**
FVP-7 REPRODUCIBILITY	**UNTESTED.**

FVP TABLE 48-D — The manufacturing model `ln r_t = ln D − ln p_bar − ln(1 + (N−1)u_0(1−λ)^t)`

Gate	Finding
FVP-1 DIMENSIONS	All terms in nats; λ dimensionless per period; t a count. **PASS.**
FVP-2 SCALE TYPE	Inherits Chapter 47's 47-A/47-C conditions plus Chapter 46's 46-F on u decay. **CONDITIONAL, and the conditions are declared upstream.**
FVP-3 LIMITING CASES	t –> inf: u –> 0, N_eff –> 1, ln r –> ln(D/p_bar) — the parallel wall, and r **converges** rather than diverging, which is the correct behaviour: over-framing has a ceiling set by the doctrine load, not an unbounded runaway. λ –> 0: r constant, no drift. **PASS.**
FVP-4 FIXED POINTS	u = 0 is globally attracting for λ > 0 (Ch 46, FVP 46-F). The crossing t∗ is not a fixed point but a threshold, and it exists iff u_0 > u∗. **PASS.**
FVP-5 IDENTIFIABILITY	λ from two clean u measurements; u_0 from the channel audit; D, p_bar, N from Chapter 47's protocols. All separately obtainable. **t∗ is then a prediction with no free parameters, which makes it directly testable. PASS.**

Gate	Finding
FVP-6 FALSIFICATION	**The strongest test in Part IX: preregister t∗ from measurements at t = 0, then measure r at t∗ and at 2t∗.** Kill condition: r fails to cross 1 within a preregistered tolerance of t∗, or the constant-u baseline predicts the observed r series as well as the decay model.
FVP-7 REPRODUCIBILITY	**UNTESTED.** Whether λ recovered in one organisation transfers to another of comparable structure is unknown.

The Specification

DD-1 . Diagnose direction before magnitude. Two instruments minimum, one behavioural and one structural, with I-4 preferred as the structural one. [direction]

DD-2 . Report ln r signed. Never report Δ bare. §48.2. [reporting]

DD-3 . Compute r per decision-making unit, never for the whole organisation. §48.11. [measurement] *Test: the report contains one ln r per unit and no organisation-wide D/P.*

DD-4 . Aggregate S, intake-weighted — never aggregate D and P. The aggregate ratio overstated bankable input by a factor of 4.26 in §48.11's case. [measurement]

DD-5 . Run the infeasibility check before designing any correction. If D/p_bar > N, only doctrine reduction works; record Δ_min = ln(D/p_bar) − ln N and say so in the report. [feasibility]

DD-6 . Derive thresholds from the organisation's own S_min = C/η, not from a literature. No empirical distribution of institutional Δ exists. [thresholds]

DD-7 . Demote σ∗- and κ∗-dependent instruments to orderings where the baseline is undefended, in print, per FVP-2 and Chapter 53. [honesty]

DD-8 . Treat high anomaly absorption as conclusive and low absorption as inconclusive. The instrument's bias runs one way (FVP 48-B/FVP-5). [I–2]

DD-9 . Compute t∗ and diarise it. t∗ = ln(u∗/u_0)/ln(1−λ), with u∗ = (D/p_bar − 1)/(N−1). Re-measure before it, not after. [drift] *Test: t∗ is written down with its date of expiry.*

DD-10 . Install the u maintenance schedule before declaring an over-framing correction complete. Otherwise t∗ simply restarts. [u, maintenance]

DD-11 . Rank doctrine for retirement by decisions-constrained per unit of supporting evidence, not by likelihood of being wrong. [0–4]

DD-12 . Correct under-framing by building the frame, not by cutting channels. Triage only channels with measured near-zero pairwise u. [U–1, U–2]

DD-13 . Require quantified forecasts with resolution dates from any unit claiming calibration. Their absence is diagnostic (I-8). [I–3, I–8]

DD-14 . Set the over-framing alarm at the moment an under-framing correction lands. Raising D shortens $t*$. §48.13, U-STEP 7. [loop closure]

DD-15 . Preregister every correction — encoder, fit window, freeze, held-out window, four baselines including the shuffled-input control, and the kill conditions. Chapter 47, G-14. [falsifiability]

DD-16 . No instrument or correction here may remove a person's ability to decline — to decline being measured, to decline a channel role, or to decline the dissent the protocol asks of them. CANON §5.4. [ethics]

The Failure Modes

FM-48.1 . THE UNSIGNED DIAGNOSIS

Description. Δ reported without the sign of $\ln r$; the correction is chosen by reflex. **Damages.** direction; the intervention worsens the condition. **Signature.** An under-framed unit receiving a diversity, breadth or collection intervention, and performing worse afterwards. **Measurement.** Read the report. Is $\ln r$ signed? **Fix.** DD-2. Re-diagnose; reverse the lever.

FM-48.2 . THE CANCELLING AGGREGATE

Description. Organisation-wide D/P computed from summed numerators and denominators, with units distorted in opposite directions. **Damages.** measurement, in the dangerous direction — bankable input overstated 4.26x in §48.11. **Signature.** A near-calibrated headline figure and universally poor unit-level performance. **Measurement.** Recompute per unit (DD-3). **Fix.** Aggregate S, intake-weighted, never D and P.

FM-48.3 . THE ONE-OFF DIVERSITY PROJECT

Description. u raised once, as a project, with no maintenance schedule. **Damages.** u, on a delay. $t*$ restarts from the new u_0. **Signature.** A correction that works, is declared complete, and is followed by the same diagnosis two to four years later. **Measurement.** Two u measurements after the intervention, separated by six or more periods. **Fix.** DD-10. The lever is a schedule, not a project (Ch 46, S-12).

FM-48.4 . THE COLLECTION REFLEX

Description. An under-framed organisation responding to its own confusion by collecting more. **Damages.** r falls further; the loop is self-reinforcing. **Signature.** Instrument proliferation with flat citation (I-7); a rising backlog; a standing request for more data that is never satisfied. **Measurement.** I-6 and I-7 together. **Fix.** U-STEP 3. Build the frame. Stop collecting against a frame that does not exist.

FM-48.5 . THE TRIAGE OVERSHOOT

Description. Under-framing corrected by cutting independent channels rather than by building doctrine. **Damages.** u, N_eff, and every future correction's ceiling. **Signature.** Calibration achieved, followed by rapid drift into over-framing — §48.10's trajectory, entered deliberately. **Measurement.** N_eff before and after; t∗ computed from the post-triage u_0. **Fix.** Restrict triage to channels with measured near-zero pairwise u (DD-12). Rebuild N_eff before it is needed.

FM-48.6 . THE INVISIBLE ANOMALY

Description. Severe over-framing suppresses anomaly recognition, so the anomaly audit returns a low absorption rate and reads as healthy. **Damages.** the diagnosis. **Signature.** Low A, low σ, and a high surprise rate from outside the model's categories. **Measurement.** Cross-check I-2 against I-1 and I-4. A low A with a low σ is the signature; a low A with a healthy σ is genuine. **Fix.** DD-8. Never read a low A alone as calibration.

FM-48.7 . THE SCOPE EXPANSION

Description. Over-framing addressed by raising N — broadening the declared domain — because it is the politically easiest lever. **Damages.** C rises; ln r barely moves; the underlying u problem is untouched. **Signature.** A remit that grows while distortion holds steady. **Measurement.** ln r before and after. Raising N alone does not move it; it only raises the ceiling on what u could later buy. **Fix.** O-1 or O-4. Chapter 8's ranking stands: N is the most expensive lever and the least rewarding.

FM-48.8 . THE CORRECTED-INTO-THE-OTHER-DITCH

Description. An under-framing correction lands, D rises, and no one recomputes t∗. **Damages.** the loop is not closed; over-framing arrives on schedule and is treated as a new problem. **Signature.** An organisation with a documented history of alternating between "we have no direction" and "we are not listening", every few years. **Measurement.** t∗ from the post-correction D, u_0, λ. **Fix.** DD-14. The alarm is set at the moment of correction, not after the next failure.

Worked Example — A Two-Site Agency

A constructed case carried end to end. **Every number is the example's own; none is drawn from a real organisation and none should be cited as empirical.** The arithmetic is reproducible from the stated parameters.

The system

An agency with two decision-making units: a headquarters that sets policy and a field operation that executes and observes. Both work the same N = 12 decision domains. Quarterly periods. Storage parameters shared: η = 0.42, δ = 0.14, C = 0.55 per unit, so S_min = C/η = 1.3095 nats and the operating envelope is min(r, 1/r) >= 1.3095/2.4849 = 0.5270, i.e. 0.527 <= r <= 1.898.

1. Baseline, per unit (DD-3)

Channel audit (I-4) and doctrine enumeration (Chapter 47, G-5), `p_bar = 1.0` in both units:

unit	u	N_eff = P	D	ln r	r	Δ	S = lnN.min(r,1/r)
HQ	0.11	2.210	6.80	+1.1239	3.0769	1.1239	0.8076
Field	0.62	7.820	1.15	−1.9169	0.1471	1.9169	0.3654

HQ is over-framed at Δ = 1.12. The field is under-framed at Δ = 1.92. Both are outside the envelope. Both would have equilibrium stores well below zero.

The aggregate, for contrast only (DD-4 forbids using it):

`D_total = 7.95` `P_total = 10.030` `r_agg = 0.7926` `Δ_agg = 0.2324`

The agency's headline figure says it is nearly calibrated. It is not; it has one unit at 3.08 and one at 0.147. Computing bankable input from the aggregate gives 1.9696 nats against a true intake-weighted 0.4629 — an overstatement of 4.26 times.

2. Corroboration

The instruments, read per §48.9:

	HQ	FIELD
I-1 revision rate	low	high
I-2 anomaly absorption A	high, conclusive	low, inconclusive (FM-48.6 check)
I-3 calibration	overconfident	no quantified forecasts
I-5 intake per decision	low	high
I-6 backlog	flat	rising
I-8 forecasts present	yes	no

Two instruments agree in each unit, one behavioural and one structural (DD-1). The field's low A is read as inconclusive per DD-8, and the FM-48.6 cross-check clears it: the field's σ is high, so its low absorption is genuine and not suppression.

3. Infeasibility check (DD-5)

```
HQ:    D/p_bar = 6.80  vs  N = 12   -> FEASIBLE.  u* = (6.80−1)/11 = 0.5273
Field: D/p_bar = 1.15  vs  N = 12   -> FEASIBLE, but the fix runs the other way.
```

HQ can be calibrated by geometry, but it needs `u = 0.5273` from a measured 0.11 — a factor of 4.8. That is a very large move and prices O-4 into the conversation immediately.

4. The obvious move, and why it is insufficient

The agency's instinct will be redistribution: move doctrine from HQ to the field. Under strict conservation of D and P:

`D_total = 7.95` `P_total = 10.030`

```
Best achievable by redistributing D alone (P fixed):
    D_HQ = 2.210, D_field = 5.740
    ->  r_HQ = 1.0000,  r_field = 0.7340,  mean Δ = 0.1546

Best achievable by redistributing both (any split):
    both units at r = D_total/P_total = 0.7926,  Δ = 0.2324 each
```

Redistribution gets both units inside the envelope. It cannot calibrate both, because the totals do not match: the agency is doctrine-poor overall by `10.030 − 7.95 = 2.080` units, 26.2 per cent of its standing doctrine.

That is a genuinely useful finding and it is invisible without per-unit measurement. The agency does not have a distribution problem alone. It has a distribution problem *and* a shortfall, and the shortfall is in doctrine, which is the opposite of what the HQ diagnosis on its own would suggest.

5. The intervention

Stage 1 — redistribute (both units into the envelope).

Move 4.59 units of doctrine from HQ to the field: HQ to `D = 2.21`, field to `D = 5.74`.

```
HQ:     r = 2.210/2.210 = 1.0000   Δ = 0.0000   S = 2.4849
Field:  r = 5.740/7.820 = 0.7340   Δ = 0.3092   S = 1.8239
```

Both inside `0.527 <= r <= 1.898`. Intake-weighted S rises from 0.4629 to `(2.210 x 2.4849 + 7.820 x 1.8239)/10.030 = 1.9695` — a factor of 4.25.

What "moving doctrine" means concretely: HQ retires standing guidance and the field adopts explicit standing assumptions and publishes forecasts. **It is not a transfer of documents. It is HQ ceasing to determine what the field will find, and the field starting to commit to what it expects.**

Stage 2 — the shortfall. The agency still needs 2.080 units of net new doctrine to bring both units to calibration. Per O-STEP 5 and U-STEP 3 this is built, not moved: explicit models, committed hypotheses, published forecasts with resolution dates, and closed lines of enquiry.

Stage 3 — HQ's u. HQ sits at `u = 0.11`, `N_eff = 2.21`, in an agency of this size. Even after Stage 1 it is calibrated at a very low effective count, which means it has almost no margin: any future growth in `D` puts it straight back over the line. Raise u toward 0.30 (`N_eff = 4.30`) as insurance, on a maintenance schedule (DD-10).

6. The alarm (DD-9, DD-14)

The field has just received doctrine. Compute its crossing date now.

Post-Stage-1 field: `D = 5.74`, `p_bar = 1.0`, `u_0 = 0.62`, and suppose `λ = 0.06` per quarter measured per Chapter 46 S-11.

```
u* = (5.74 − 1)/11 = 0.4309
t* = ln(0.4309/0.62) / ln(0.94) = 5.88 quarters
```

The field crosses from under-framed to over-framed in under six quarters, and the agency should diarise re-measurement at quarter 4, not quarter 8. Without the alarm, the field's correction becomes the field's next pathology, and the agency acquires the alternating history of FM-48.8.

7. What would falsify this

Preregistered per DD-15 and FVP 48-D/FVP-6:

1. `ln r` measured per unit quarterly for eight quarters after the intervention.
2. Held-out prediction of unit-level decision quality, assessed at 24-month outcome observation, for quarters 9-16.
3. Four baselines: persistence; `E` alone; matched autoregressive; and the shuffled-`ln r` control.
4. **The specific, cheap, decisive test:** the field's predicted crossing at $t* = 5.88$ quarters. $t*$ is computed from three independently measured quantities with **no free parameters**. Kill condition: the field's `r` fails to cross 1 within a preregistered tolerance around quarter 6, or a constant-`u` model predicts the observed `r` series as well as the decay model.

Point 4 is the one to run first. It is the strongest falsification opportunity in Part IX because it is a dated, parameter-free prediction about an observable quantity, and it can be checked with a channel audit and a decision log.

Both pathologies are real, both are common, both are measurable from records organisations already keep, and their corrections are opposites. The over-framed institution must lose doctrine or gain independence; the under-framed institution must gain doctrine and must not gain independence. **And the mechanism that carries an organisation from one to the other is not a decision anyone makes — it is the decay of `u`, which is Chapter 11's pseudo-diversity operating as a clock.** An organisation that measures its distortion once has taken a photograph. An organisation that computes $t*$ has read the timetable.

Chapter 49 . Trust, Obligation, and Reciprocity as Quantities

Three words do more rhetorical work in the literature of institutions than any others, and all three are usually left as adjectives. A team is described as *high-trust*. A community is said to have *strong reciprocal norms*. A network is *dense with obligation*. None of those sentences tells an engineer what to build, what to measure, or what would falsify the claim.

This chapter converts all three into quantities that enter the storage law at a named position, with a stated unit, an obtainable measurement, and a computed price. The conversion is not decorative. Each of the three attaches to a *different* parameter, which is why they behave differently, why they fail differently, and why the intervention that fixes one will not fix another.

The three claims this chapter develops, stated at the front so that the rest can be checked against them:

- **Trust is a delta-reduction device.** It lowers the retention loss on stored order, because a trusted agreement needs less re-verification per period to remain in force. Trust does not add anything to the input side of the law. It reduces the rate at which what you already have leaks away, and it does so by making a specific expenditure — verification — unnecessary.
- **Reciprocity is an eta-raiser.** It converts interaction into durable relation. Interaction that is not returned deposits nothing that both parties hold; interaction that is returned creates a two-sided claim, and a two-sided claim is structure.
- **Obligation is the antisymmetric part of a relation.** Relationship volume is the symmetric part; obligation is the signed remainder. This is not an analogy — it is a decomposition of the interaction matrix, and both parts are computable from the same data.

Chapter 46 built the dyad. Chapters 47 and 48 built the group's moment and its doctrine-to-data fault. This chapter supplies the three quantities that determine whether anything built in those chapters survives a year.

The Law

49.0.1 What is proven

Nothing in this chapter adds to Papers A through E. What is proven, and quoted here unaltered:

```
F(u) = ln(1 + (N−1)u)                          Paper A, Theorem 2.1
N_eff = 1 + (N−1)u                             effective count
0 <= F(u) <= ln N                              Paper D, Corollary 2.2

Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t        Paper E, Section 3
S_t = F(U_t),  0 <= S_t <= ln N

admissibility:  C_t <= (1−delta)Q_t + eta S_t   guarantees Q_{t+1} >= 0
one−step growth: Q_{t+1} > Q_t  <=>  eta S_t > delta Q_t + C_t     Cor 3.2
steady state:    Q_inf = (eta F(u∗) − c∗) / delta                   Thm 3.5
headroom:       h_t = eta ln N − delta Q_t − C_t                 Section 6.4
```

Theorem 3.5 carries a clause that this chapter leans on twice: **if `eta F(u∗) < c∗` the equilibrium is negative, and the theorem does not hide it** — the regime is unsustainable, or the linear model

must be replaced. Section 49.10 shows a real case where a trust intervention makes a negative equilibrium *more* negative, and the clause is the reason.

49.0.2 What is postulated, badged

Three mappings are introduced below. **Each is a postulate — a modelling choice about how a measured social quantity enters a proven equation — and none of them is derived from Papers A–E.** They are stated here with their badges on so that the badge does not have to be re-attached at every use.

Postulate T (trust enters as delta). For a relation in which stored order consists of standing commitments, the per-period retention loss decomposes as

```
delta_eff = delta_0 + (1 − a)(1 − v)
```

where a in [0,1] is the **reliance rate** — the measured probability that an unverified commitment is honoured within the period — v in [0,1] is the fraction of commitments actually re-verified, and `delta_0 >= 0` is the verification-irreducible loss (death, departure, genuine forgetting, external shock). The postulate is that verification is *effective*: a verified commitment does not lapse for want of confirmation.

Postulate R (reciprocity enters as eta). For a relation carrying directed interaction volumes `x_ij` and `x_ji` over an accounting window, only the *matched* portion of the volume converts into durable structure:

```
eta = eta_max ∗ rho,     rho = 2 min(x_ij, x_ji) / (x_ij + x_ji)
```

where `eta_max` is the conversion efficiency the relation would have at perfect balance. The postulate is the word *only*. The arithmetic of `rho` is a definition, not a claim.

Postulate O (obligation is the antisymmetric part). Writing the interaction matrix X as its symmetric and antisymmetric parts,

```
X = Sym + Skew,    Sym_ij = (x_ij + x_ji)/2,    Skew_ij = (x_ij − x_ji)/2
```

`Sym` measures the relation and `Skew` measures the obligation. The decomposition is exact and unique for any square matrix; the postulate is the *interpretation* of the two parts.

49.0.3 Honesty clauses

Repeated because this chapter is the part of Part IX most likely to be quoted in a report.

- As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. Every price quoted below is a floor.
- The triangle inequality **fails** for geometric triples (Paper A §5). Subadditivity holds on the parameter u only. Three-party obligation chains cannot be composed metrically.
- Paper A is pure mathematics and makes **no empirical or metaphysical claim**. Surplus is not identified with Shannon entropy or mutual information (Paper D). **No thermodynamic negentropy law is proved** (Paper E).
- Gambetta, Putnam and Ostrom are cited below for what they actually reported. **The transfer of their findings into this parameterisation is a modelling assumption, badged at each use.**

- CANON §5.4 governs absolutely: **no intervention in this chapter may remove a person's ability to decline.** Section 49.11 shows that this is not a decoration but a load-bearing constraint, because the highest-a obligation networks on record are precisely the ones that violate it.

The Engineering

49.1 Trust, defined so that it can be measured

Start by refusing the usual definition. Trust is commonly defined as a belief or a feeling — an expectation held by one party about another. That definition is unmeasurable except by asking, and asking produces a self-report on an ordinal scale, which by FVP-2 cannot be divided, multiplied, or entered into a ratio.

The operational definition this manual uses instead:

> **Trust on a relation is the measured rate at which unverified commitments are honoured.**

Call it a, the **reliance rate**. It is a probability, it is dimensionless, it has a true zero (no unverified commitment ever holds) and a true ceiling (all of them do), and it is estimable from records that most organisations already keep. It is a property of a *relation in a domain*, not of a person: your supplier's a on delivery dates and their a on cost estimates are different numbers and must be measured separately.

The point of the definition is the word "unverified." A commitment you checked does not tell you anything about trust; it tells you about your checking. Trust is only visible in the commitments you did not check, which is why measuring it requires deliberately not checking some of them, and why that in turn requires an ethics protocol (§49.15, Protocol A).

49.2 Why trust reduces delta and not anything else

Consider a relation whose stored order `Q` is a set of standing commitments: agreements about who does what, who may draw on what, what a given signal means. This is the ordinary content of institutional `Q`.

Each period, a commitment either carries forward or lapses. Lapse has two causes:

1. **Irreducible causes.** The person leaves, the memory genuinely goes, the world changes under the agreement. No amount of confirmation prevents these. Rate: `delta_0`.
2. **Confirmation-addressable causes.** The commitment is still nominally in force but nobody has confirmed it, and in the absence of confirmation the parties diverge on what it says or whether it binds. Rate, per unconfirmed commitment: `1 − a`.

If a fraction v of commitments is re-verified each period, the confirmation-addressable loss applies only to the remaining `1 − v`:

```
delta_eff = delta_0 + (1 − a)(1 − v)
```

That is Postulate T, and its content is immediate: **there are two ways to hold delta down, and trust is the one that does not cost anything per period.** Raising a and raising v reduce the

same term. Raising v costs `C_v = c_v * v` in the dissipation budget. Raising a costs nothing recurrent.

This is the precise sense in which **trust is a delta-reduction device**: it occupies the exact position in the storage law that verification spending occupies, and it is free at the margin.

49.3 The frontier, and what trust does to it

Fix a. Varying v from 0 to 1 traces a straight line in the (`delta`, `C_v`) plane:

```
delta = delta_0 + (1 − a)(1 − C_v / c_v)
```

Endpoints, computed with `delta_0 = 0.04` and `c_v = 0.25`:

a	v = 0 (verify nothing)	v = 1 (verify everything)
0.50	delta = 0.5400, C_v = 0.0000	delta = 0.0400, C_v = 0.2500
0.70	delta = 0.3400, C_v = 0.0000	delta = 0.0400, C_v = 0.2500
0.85	delta = 0.1900, C_v = 0.0000	delta = 0.0400, C_v = 0.2500
0.95	delta = 0.0900, C_v = 0.0000	delta = 0.0400, C_v = 0.2500

Read the table by its intercepts. **Trust moves the delta-intercept toward the origin and leaves the C-intercept exactly where it was.** Every point on a high-a frontier weakly dominates the corresponding point on a low-a frontier: same verification spend, lower loss. That is the geometric statement of the dividend, and it is why trust is not a substitute for verification but a *rotation of the whole trade-off*.

The half-life consequence, with `v = 0`, computed as `t_half = ln 2 / −ln(1 − delta)`:

a	delta_eff	half-life of an unmaintained agreement (periods)
0.50	0.5400	0.893
0.70	0.3400	1.668
0.80	0.2400	2.526
0.90	0.1400	4.596
0.95	0.0900	7.350
0.99	0.0500	13.513

An agreement in a relation at `a = 0.50` is half-forgotten before the next period ends. The same agreement at `a = 0.99` survives thirteen and a half periods. **The agreement is identical. The paper is identical. What differs is the rate at which the paper stops describing what people do.**

49.4 The steady-state dividend, and why it is superlinear

Steady state, Paper E Theorem 3.5, with `N = 8`, `u = 0.55`, so `S = ln(1 + 7(0.55)) = 1.5790`; `eta = 0.35`; `C_0 = 0.30`; hence the net input `K = eta*S − C_0 = 0.2526`:

```
Q_inf = K / delta_eff = 0.2526 / (0.04 + (1−a))
```

a	delta_eff	Q_inf	share of the a = 0.99 ceiling
0.50	0.5400	0.4679	9.3%
0.60	0.4400	0.5742	11.4%
0.70	0.3400	0.7431	14.7%
0.80	0.2400	1.0527	20.8%
0.85	0.1900	1.3297	26.3%
0.90	0.1400	1.8046	35.7%
0.95	0.0900	2.8071	55.6%
0.97	0.0700	3.6092	71.4%
0.99	0.0500	5.0529	100.0%

Because Q_inf varies as 1/delta and delta varies as 1 − a, the elasticity of stored order with respect to the reliance rate is a/(1−a):

a	0.50	0.70	0.80	0.90	0.95	0.99
elasticity of Q_inf in a	1.00	2.33	4.00	9.00	19.00	99.00

A one per cent improvement in the reliance rate at a = 0.95 buys nineteen per cent more stored order; the same one per cent at a = 0.50 buys one per cent. This is the arithmetic behind an observation that is otherwise merely folkloric: the last increments of reliability are worth absurdly more than the first, and organisations that are nearly reliable are qualitatively different from organisations that are mostly reliable. It also explains the reverse — why a single prominent breach in a high-a institution is so destructive. It is not sentiment. It is 1/(1−a).

The honest caveat, stated in the same breath: this superlinearity is an artefact of the model's 1/delta steady state, which assumes the linear law holds all the way to delta -> 0. It does not. As delta approaches delta_0 the irreducible term dominates and the curve flattens; the table above already includes delta_0 = 0.04, which is why the a = 0.99 entry is 5.05 and not 25.26. Anyone quoting the elasticity must quote delta_0 with it.

49.5 When to verify, and the result that verification is all-or-nothing

Take Q_inf(v) = (K − c_v v)/(delta_0 + (1−a)(1−v)) and differentiate. Writing A = 1 − a:

```
sign( dQ_inf/dv ) = sign[ A (K − c_v) − c_v ∗ delta_0 ]
```

The sign does not depend on v. **The objective is a linear-fractional function of the verification rate, so its optimum is always at an endpoint: verify everything, or verify nothing.** The switch happens at

```
a∗ = 1 − c_v ∗ delta_0 / (K − c_v)        valid when K > c_v
```

With K = 0.2526, delta_0 = 0.04:

unit verification cost c_v	a*	reading
0.02	0.9966	verify unless the relation is essentially perfect
0.06	0.9875	verify unless a exceeds 98.75%
0.12	0.9638	verify below 96.4%
0.20	0.8479	verify below 84.8%
0.24	0.2381	verify only in near-total distrust
>= 0.2526	never	verification consumes the entire net input

Check the two ends against intuition, because a formula that surprises you at the ends is usually wrong. **Cheap verification means you should verify almost always** — which is why retail commerce, where verification is a barcode scan, does not run on trust and does not need to. **Expensive verification means you should verify almost never** — which is why relationships whose content is judgement, taste, discretion or intent run on trust and *have no alternative*, since auditing a judgement costs more than the judgement is worth.

This is the whole practical content of the trust parameter: trust is valuable exactly in proportion to the cost of the verification it displaces. Where verification is cheap, trust is a luxury. Where verification is expensive or impossible, trust is the only available delta control, and its absence is not a cultural deficiency but a hard ceiling on stored order.

Model artefact, flagged. The bang-bang result follows from the *linearity* of `C_v = c_v v`. Real verification costs are convex — the first ten per cent of checks are cheap and the last ten per cent are punishing — and with a convex cost the optimum is interior. Use the endpoint result as a signpost about *which direction* to move, never as an instruction to audit everything or nothing. What survives convexification is the threshold structure: there is an a above which additional verification destroys more order than it saves.

49.6 What trust is not

Four negations, each of which is a common and expensive error.

Trust is not scalar. There is no such thing as an organisation's trust level. There is a matrix `a_ij^(d)` indexed by pair and by domain. Averaging it discards the only information it carries. A supplier at `a = 0.98` on delivery and `a = 0.55` on cost forecasts is not an `a = 0.77` supplier; they are a supplier you rely on for one thing and audit for another, and the average would have you doing exactly the wrong thing in both domains.

Trust is not a feeling. The feeling is a predictor of a, sometimes a good one and sometimes catastrophically bad one, and it is available before the data. Treat it as a prior, not as the measurement. Where prior and measurement disagree, the measurement is the datum.

Trust is not an input. It appears nowhere in `eta S_t`. A high-trust group with no interaction surplus accumulates nothing; it merely loses nothing slowly. Chapter 46's canon holds: **the wedding of like with like is a non-event**, and it remains a non-event however much the parties rely on each other. Trust protects the store; it does not fill it.

Trust is not always desirable. `delta` is not a fault to be minimised (Chapter 3). A system with no forgetting cannot learn. High a on a commitment that ought to lapse is exactly the mechanism by which obsolete procedure outlives its purpose — the commitment is honoured, unverified, forever,

because nobody re-opens what is trusted. **The engineering target is not maximum a; it is a high on the commitments that should persist and deliberately low on the ones that should not.** That is what a sunset clause is: a manufactured reduction in a on a specific agreement, scheduled in advance.

49.7 Reciprocity, defined so that it can be measured

Reciprocity is usually described as a norm — the expectation that a favour will be returned. As a norm it is unmeasurable; as a *flow balance* it is trivially measurable, provided you can count the flows.

Let `x_ij` be the volume of transfers from `i` to `j` over an accounting window `w`, in some common unit. The **matched volume** of the dyad is the part that both parties supplied:

```
matched_ij = 2 min(x_ij, x_ji)
total_ij   = x_ij + x_ji
rho_ij     = matched_ij / total_ij = 2 min(x_ij, x_ji) / (x_ij + x_ji)
```

Equivalently, and this identity is worth carrying because it is the bridge to obligation:

```
rho_ij = 1 − |x_ij − x_ji| / (x_ij + x_ji)
```

The two expressions are the same number, since `a + b − |a − b| = 2 min(a,b)`. **The reciprocity index is simultaneously the matched fraction of the volume and one minus the normalised imbalance.** It is dimensionless, has a true zero (all flow one way) and a true one (perfect balance), and is invariant to the unit in which flows are counted — all three properties are required by FVP-1 and FVP-2 before it can be multiplied by anything.

x_ij	x_ji	matched	total	rho	eta = 0.45 ∗ rho
10	10	20	20	1.0000	0.4500
10	8	16	18	0.8889	0.4000
10	6	12	16	0.7500	0.3375
10	5	10	15	0.6667	0.3000
10	4	8	14	0.5714	0.2571
10	2	4	12	0.3333	0.1500
10	1	2	11	0.1818	0.0818
10	0	0	10	0.0000	0.0000

49.8 Why reciprocity raises eta

`eta` is the fraction of interaction surplus that becomes durable structure (Chapter 4). Postulate R asserts that the durable fraction is the matched fraction, and the reasoning is this:

An unreturned transfer creates a claim held by one party. A claim held by one party persists at the *debtor's* convenience; it is an entry in one ledger, and the other ledger has no matching row. There is no mechanism by which both parties are separately motivated to maintain it, and there is no independent copy against which drift can be detected. A returned transfer creates a claim held by both. Both maintain it, both can check it, and it survives the departure of either one's memory of the original transaction.

Structure that only one party holds is not structure; it is a debt. That sentence is the postulate, and it is the reason `eta = eta_max * rho` rather than `eta = eta_max`.

The consequence at steady state is severe, because `eta` sits in the numerator while `delta` sits in the denominator. With `S = 1.5790`, `C = 0.30`, `delta = 0.10`, `eta_max = 0.45`:

x	y	rho	eta	eta S	Q_inf = (eta S − C)/delta
10	10	1.0000	0.4500	0.7106	+4.1055
9	6	0.8000	0.3600	0.5684	+2.6844
8	3	0.5455	0.2455	0.3876	+0.8757
12	2	0.2857	0.1286	0.2030	-0.9699
5	0	0.0000	0.0000	0.0000	-3.0000

The negative entries are not small numbers; they are the theorem's own statement that the regime is unsustainable (Thm 3.5) and that the linear model must be replaced. **A relation running at `rho = 0.29` does not accumulate slowly. It fails to accumulate at all, and it does so while the interaction volume — fourteen transfers a period — looks perfectly healthy on a dashboard.**

49.9 The window problem, stated honestly

`rho` is not a property of a relation. It is a property of a relation *and an accounting window*. A favour returned in month nine counts as reciprocated if the window is a year and unreciprocated if the window is a quarter. Fifteen favours, with the following return lags in months, five of which were long and two of which never returned:

window w (months)	1	2	3	6	12	24
estimated `rho`	0.1333	0.2667	0.4000	0.6000	0.8000	0.8667

The same relation measures at `rho = 0.13` and `rho = 0.87` depending on a choice the analyst makes. This is a genuine identifiability problem (FVP-5) and there are only three honest responses:

1. **Preregister the window** before looking at the data, and justify it from the domain's natural cycle — the length of a project, a season, a budget year.
2. **Report the whole curve**, as above, not a single number. The *shape* is informative: a curve that rises steeply and plateaus early describes a fast-settling relation; one that keeps rising at 24 months describes a relation carrying long debt.
3. **Set the window from the decay rate.** The natural choice is `w = 1/delta`, the time constant: beyond that horizon the structure the return was supposed to build has already leaked away, so a return arriving later is arriving into an emptier store. This is the principled option and the one this manual recommends.

49.10 Obligation as the antisymmetric part

Now the third quantity. Take the interaction matrix X and split it:

```
Sym_ij  = (x_ij + x_ji)/2     the relation:   how much passes between i and j
Skew_ij = (x_ij − x_ji)/2     the obligation: which way it is out of balance
```

Every square matrix has exactly one such decomposition. Sym is the volume Chapter 46 called the dyad's traffic; Skew is the standing debt. Note immediately that `rho_ij = 1 − |Skew_ij|/Sym_ij`, so reciprocity and obligation are two readings of one measurement, not two measurements.

At the node level, define each member's **net obligation balance** as the row sum of Skew:

```
b_i = sum_j (x_ij − x_ji) = (total given) − (total received)
```

`b_i > 0` means i is a net creditor. Necessarily `sum_i b_i = 0`. Define the **network imbalance index**:

```
B = sum_i |b_i| / (2 ∗ total volume)
```

`B = 0` is perfect balance at every node; B approaching 1 is a pure one-directional flow.

A small theorem worth stating, because it prevents a common error. Node-level balance can hide dyad-level imbalance, since a member who is a creditor to one partner and a debtor to another shows `b_i = 0`. Formally `|sum_j (x_ij − x_ji)| <= sum_j |x_ij − x_ji|`, so B never exceeds the dyad-level imbalance, and therefore

```
1 − B  >=  rho
```

`1 − B` is an upper bound on reciprocity, never an estimate of it. In the worked consortium at the end of this chapter the two numbers are `1 − B = 0.4412` and `rho = 0.3824`: the node-level index overstates the health of the network by fifteen per cent, and it does so systematically, in the direction of comfort. Report both or report `rho`.

49.11 Obligation networks: Gambetta and Putnam, read quantitatively

Two of the most cited empirical bodies on this subject describe the same structural object from opposite ends, and the parameterisation above separates them cleanly.

Gambetta (1993) studied the Sicilian case: an order maintained through networked obligation in a setting where central enforcement had failed. Read in these terms, the finding is that a protection market arises precisely where a is low and c_v is high — where commitments are not honoured unverified, and verifying them independently is prohibitively expensive. What is sold in such a market is not a good but a **delta reduction**: a third party guarantees that the commitment holds, so the parties need not verify. That is an economically coherent product, and its existence is exactly what the frontier of §49.3 predicts when a is low and c_v is unaffordable.

And here the chapter must state the boundary rather than admire the mechanism. The guarantee is enforced by removing the ability to decline. Under CANON §5.4 that is not a permitted intervention, and under the framework's own accounting it is not even a good one: coerced compliance raises *measured* a while leaving the underlying relation at `rho` near zero, because the flows are one-directional by construction. The system therefore shows a low `delta` and a `eta`

near zero simultaneously — which by Theorem 3.5 is a negative equilibrium propped up by continuous coercive input. **An obligation network that removes exit is not a trust system. It is a capture system, and the two are distinguishable by measuring `rho`, which the first has and the second does not.** That is a testable difference, and it is the reason the ethical rule and the engineering rule point the same way here.

Putnam (1993) studied the divergence in institutional performance between northern and southern Italian regions after the 1970 devolution of identical formal structures, and attributed the difference to accumulated civic tradition — *social capital*. The mapping is direct and it is the cleanest available illustration of the whole storage law at civilisational scale:

```
social capital      <-> Q, stored order held on relations rather than in a centre
civic association    <-> the input term eta S: interaction that leaves structure
horizontal networks  <-> high rho: mutual, balanced, both parties hold the claim
vertical patron-client<-> low rho: one-directional, only the patron holds the claim
tradition's persistence<-> low delta: Q with a very long time constant
```

Putnam's most striking claim — that the divergence traces to civic patterns centuries old — is, in this notation, a claim that `delta` is very small: a store with a half-life measured in generations. **This is exactly the regime in which `Q_inf = (eta S - C)/delta` becomes enormous for even modest input, and it is also the regime in which recovery from a low `Q` is hopelessly slow, because the same small `delta` that preserves a good store preserves a bad one.** Low delta is symmetric. It is not a virtue; it is a memory.

The honest note, required. Putnam's causal argument has been contested by historians and political scientists on the strength of the historical evidence and on the direction of causation — whether civic tradition produced good government or good government produced civic tradition. This manual takes no position on that dispute. What the framework offers is not a verdict but a decomposition: the two hypotheses make *different predictions about which parameter moves first*, and that is a measurable difference rather than an interpretive one.

49.12 Ostrom's design principles, from the quantitative side

Ostrom (*Governing the Commons*, 1990) derived eight design principles common to long-enduring common-pool resource institutions. Chapter 4 used them as eta-raising devices. Here each is mapped to the parameter it actually improves, which turns out to be a more discriminating exercise than the blanket attribution suggests — **they do not all improve the same thing, and three of them improve something other than `eta`.**

#	Ostrom's principle	Parameter improved	Mechanism, in this notation
1	Clearly defined boundaries	`N`, and `C`	Makes `N` a countable set rather than a rhetorical one; excludes uncounted draws that appear in the model as unexplained `C`

#	Ostrom's principle	Parameter improved	Mechanism, in this notation
2	Congruence: rules match local conditions; benefits proportional to costs	`rho`, hence `eta`	"Proportional" is literally `rho -> 1`. A rule that allocates benefit in proportion to contribution *is* a balance condition on the flow matrix
3	Collective-choice arrangements	`u`, hence `S`	Brings members' non-parallel framings into rule-making. Rules written by one framing have `u = 0` at the point where variety was needed
4	Monitoring, by monitors accountable to the appropriators	`c_v`, **not a**	See below. This is the principle most often misread
5	Graduated sanctions	preserves `rho` and `N`	A maximal first sanction deletes the edge; a graduated one preserves it. See the computation below
6	Conflict-resolution mechanisms, rapid and low-cost	`C`, directly	Unresolved conflict is a standing dissipation term. "Rapid and low-cost" is a specification on `C`, not on justice
7	Minimal recognition of the right to organise	protects `Q` from external shock	Guarantees that `Q` is not zeroed by an outside authority; in the model, it bounds a catastrophic `delta` event at 1
8	Nested enterprises	`C`, by orders of magnitude	The only principle that changes the *scaling law* rather than a coefficient. See §49.13

Principle 4 deserves its own paragraph because the usual reading is wrong. Monitoring is routinely described as "building trust." In this parameterisation it does no such thing. Monitoring by appropriators themselves — Ostrom's specific condition, that the monitors are the users or

are accountable to them — reduces c_v, the unit cost of verification, because observation is a by-product of use rather than a separate expedition. Lowering c_v moves a∗:

c_v	a∗	policy
0.20	0.8479	verify only below 85% reliance
0.12	0.9638	verify below 96%
0.06	0.9875	verify below 98.75%
0.02	0.9966	verify below 99.7%

Cheap monitoring raises the threshold above which you should stop verifying; it does not raise the reliance rate. The confusion matters because it produces the wrong intervention: an organisation that believes monitoring builds trust responds to a breach by monitoring more, which lowers delta through v at a cost in C, and leaves a exactly where it was. Ostrom's own condition — *monitors accountable to the appropriators* — is what keeps the monitoring from appearing on the C line as pure overhead, and it is the part usually dropped in transfer.

Principle 5, computed. In the eight-firm consortium of the Worked Example (post-intervention state), firm 3 is the largest net giver and a hypothetical rule-breaker. Two responses:

response	rho after	N after	S after	eta∗S	change
intact	0.8947	8	1.5347	0.6179	—
expulsion (maximal sanction)	0.8947	7	1.4159	0.5701	**-7.7%**
30% flow reduction, edge kept (graduated)	0.9038	8	1.5347	0.6242	**+1.0%**

Expulsion is priced entirely in N, and it is charged even when the expelled member was reciprocating perfectly. The consortium loses 7.7 per cent of its input by removing one of eight distinct participants, and it recovers none of that by having been right. The graduated response costs volume for two periods and returns a slightly *higher* eta∗S, because the reduction fell on the most imbalanced flows. Ostrom's fifth principle is not leniency. It is the observation that the sanction which destroys the relation charges the sanctioner the whole of ln N − ln(N−1).

49.13 Principle 8 computed: nesting changes the exponent

Brooks (1975) counted communication paths in a group of n as n(n−1)/2. Chapter 5 established the crossover: **cost grows quadratically while surplus grows logarithmically**, so for any positive per-path cost there is an n beyond which no arrangement helps. With eta = 0.40 and per-path cost c_1 = 0.0009:

n	paths	C = c_1 * paths	eta ln n	h at delta Q = 0
12	66	0.0594	0.9940	+0.9346
20	190	0.1710	1.1983	+1.0273
30	435	0.3915	1.3605	+0.9690
40	780	0.7020	1.4756	+0.7736
50	1225	1.1025	1.5648	+0.4623
60	1770	1.5930	1.6377	+0.0447
70	2415	2.1735	1.6994	**-0.4741**

Crossover at n = 61. And the crossover is exquisitely sensitive to the per-path cost:

c_1	0.0005	0.0010	0.0020	0.0050	0.0100
crossover n	85	58	39	23	16

Now nest. Ostrom's eighth principle says that large systems appearing in her sample were organised as *nested enterprises*: small units, whose delegates form the next unit, and so on. Model it as span g and depth d, giving N = g^d members and (g^d − 1)/(g − 1) groups, each carrying g(g−1)/2 internal paths:

span g	depth d	N	groups	nested paths	flat paths	ratio
3	3	27	13	39	351	1 : 9
3	4	81	40	120	3,240	1 : 27
4	3	64	21	126	2,016	1 : 16
5	3	125	31	310	7,750	1 : 25
6	3	216	43	645	23,220	1 : 36
8	3	512	73	2,044	130,816	1 : 64
10	4	10,000	1,111	49,995	49,995,000	1 : 1,000

Flat coordination cost is O(N^2). Nested coordination cost is O(N). That is not a coefficient improvement, it is a change of exponent, and it is the only structural move in the entire parameter set that has that property. The surplus ceiling eta ln N is untouched by nesting — N is still N — so nesting is a pure reduction of C at constant ceiling.

Paths per member, holding N at 1000 and varying the span:

span g	3	4	5	6	8	10	12
paths per member	1.50	2.00	2.50	3.00	4.00	5.00	6.00

Coordination cost per member converges to g/2 and therefore falls monotonically as the span narrows. **This does not mean spans of three are optimal**, and the model says why: narrowing the span increases depth, and depth is a delay and a distortion term that appears in D/P (Chapter 48), not in C. Each tier is an opportunity for doctrine to be reasserted and data to be summarised

away. The nesting law prices the coordination side only, and Alexander's semilattice argument — *A City is Not a Tree* (1965) — is the standing warning that a strict tree, which is what g^d describes, is the wrong shape for exactly this reason. **Overlapping nests, where a unit belongs to more than one parent, cost more paths than a tree and preserve the cross-links that a tree amputates.** The computation above is the tree's price list, and it is the floor, not the recommendation.

49.14 The interaction: why trust alone cannot rescue a low-reciprocity regime

This is the most useful result in the chapter and it falls straight out of Theorem 3.5.

Take a relation with `eta S < C`. The steady state is `Q_inf = (eta S − C)/delta`, and the numerator is negative. **Lowering `delta` — that is, raising trust — makes `Q_inf` more negative, not less.** The model is telling you, in the plainest available language, that a relation which is losing order faster than it makes it cannot be repaired by making it lose order more slowly. It can only be repaired by raising `eta S` or lowering `C`.

Concretely, from the worked consortium below. Before: `rho = 0.3824`, `a = 0.62`, so `eta = 0.1721`, `delta = 0.4300`, `eta S = 0.2641`, `C = 0.30`, `Q_inf = −0.0836`. After the full intervention: `rho = 0.8947`, `a = 0.88`, `eta = 0.4026`, `delta = 0.1700`, `Q_inf = +1.8701`.

Decompose the gain of `+1.9537` by moving one parameter at a time:

move	`Q_inf`	gain over baseline
baseline (both before)	-0.0836	—
reciprocity alone (`rho` to 0.8947, `a` still 0.62)	+0.7394	**+0.8229**
trust alone (`a` to 0.88, `rho` still 0.3824)	-0.2114	**-0.1278**
both	+1.8701	+1.9537
interaction term		+1.2586

Trust alone made it worse. Not marginally — it moved the equilibrium further into the region where the model declares the regime unsustainable, because at `rho = 0.38` the numerator was negative and dividing a negative number by a smaller denominator produces a larger negative number. And the interaction term, `+1.2586`, is larger than either main effect: **the two interventions are strongly complementary and neither is close to sufficient alone.**

Read that as an operating rule. **Fix the numerator first.** Reciprocity, conversion, cost — the input side — before trust. A trust programme run in an organisation that is not converting its interactions into anything durable will produce a slower, more polite decline, and it will produce survey results that say the intervention worked.

49.15 Measurement protocols

Three protocols, one per quantity. Each is written to be executable and each states its ethics constraint, because all three involve observing people.

Protocol A — the reliance rate a.

What it estimates: the probability that an unverified commitment in domain `d` on relation `(i,j)` is honoured within one period.

1. **Enumerate the commitment population.** List every commitment in domain `d` made on the relation in the last `k` periods. A commitment must have a date, a content, and a party. If your records cannot produce this list, that fact is itself the first finding: **an organisation that cannot enumerate its commitments cannot measure its trust, and its stored order is smaller than it believes.**
2. **Stratify by consequence.** Partition into low-, medium- and high-consequence tiers using a pre-declared rule.
3. **Randomise verification within the low tier only.** Select a random subset of the low-consequence commitments and do not verify them for one period. Verify all others as normal.
4. **Score.** `a_hat = (honoured, unverified) / (total, unverified)`. Compute a Wilson interval; with fewer than about 30 unverified commitments the interval will be wider than the effect you are looking for, and the honest report is the interval.
5. **Repeat per domain.** Do not pool across domains.

Ethics constraint, binding. Randomised non-verification is only permitted in the tier where non-verification causes no harm that the party has not already accepted, and every party must be told in advance that verification is sampled rather than universal. **Constructing a hidden test of someone's reliability is a covert manipulation and is prohibited under CANON §5.4** — the party cannot decline what they are not told about. The disclosure does not spoil the measurement: what is concealed is *which* commitments are unverified, not that sampling occurs, and the standard audit literature relies on exactly that distinction.

Common failure: estimating a from complaints. Complaints measure the honoured rate of high-consequence commitments among parties who complain. That is three selection effects stacked.

Protocol B — the reciprocity index `rho`.

What it estimates: the matched fraction of interaction volume on a relation, over a declared window.

1. **Declare the unit.** Flows must be counted in a unit that is at least ratio-scaled with a true zero (FVP-2) and that is *symmetric* — the same unit applies in both directions. Hours, items, referrals, review comments, defined favour-units. **If your unit is not symmetric — if you count "advice given" one way and "money given" the other — you do not have a ratio and `rho` is not defined.** Demote to a paired ordinal comparison and say so.
2. **Declare the window before looking.** Default `w = 1/delta` periods (§49.9).
3. **Build the matrix.** `x_ij` for every ordered pair.
4. **Compute** `rho_ij` per dyad, and `rho_group = sum matched / sum total` over all ordered pairs.
5. **Report the window sensitivity curve**, not only the point estimate.
6. **Report density alongside.** A network with three edges can have `rho = 1`. Reciprocity of a sparse network is a statement about the few edges that exist.

Common failure: computing `rho` from the node balances `b_i` instead of the dyad flows. That

gives `1 − B`, which by §49.10 is an upper bound and is biased toward good news.

Protocol C — the obligation ledger.

What it estimates: the standing directed debt between every pair, and the network's imbalance.

1. **Same matrix as Protocol B.** One data collection serves both.
2. **Compute** `Skew_ij = (x_ij − x_ji)/2`, `b_i = sum_j (x_ij − x_ji)`, `B = sum|b_i| / (2 ∗ total)`.
3. **Rank the dyads by `|Skew|`,** not by `Sym`. The largest relationships are not the most imbalanced ones, and it is imbalance that predicts breakdown.
4. **Check the sign structure.** A network with a small number of large net creditors and many small net debtors is a patron-client structure regardless of what it calls itself. A network with `b_i` near zero everywhere and high `Sym` is a horizontal association. These are Putnam's two cases and they are distinguishable in one afternoon of arithmetic.
5. **Do not publish individual balances without consent.** A ledger of who owes whom, made public inside a group, is a coercive instrument: it converts a private accounting into a public pressure, and it removes the ability to decline the accounting. Publish the *distribution* and the dyad-level findings to the dyad.

The Specification

Numbered, checkable requirements. A system claiming to have engineered trust, obligation or reciprocity must satisfy all of these or state which it fails and why.

S49.1 — Domain-indexed reliance. a shall be measured and reported per relation and per domain. A single organisation-wide trust figure shall not be published. Any report quoting one number shall state the domain it applies to.

S49.2 — Verification policy stated against a∗. For each domain, the unit verification cost `c_v`, the irreducible loss `delta_0`, and the net input `K = eta S − C_0` shall be estimated, a∗ computed, and the verification rate set on the correct side of it. A verification policy that has never been priced against a∗ is an unpriced recurring charge on `C`.

S49.3 — Sunset by default. Every commitment entering `Q` shall carry a review date. Absent a review date, the commitment inherits the relation's a, which for a well-run relation is high, and **an obsolete commitment in a high-a relation is immortal.** The sunset clause is the mechanism by which `delta` is restored deliberately on the subset of `Q` that should decay.

S49.4 — Reciprocity window declared in advance. `w` shall be preregistered before flows are tabulated, with its justification. Post-hoc window selection is a researcher degree of freedom capable of moving `rho` from 0.13 to 0.87 on identical data (§49.9) and shall be treated as such.

S49.5 — Symmetric unit. The flow unit shall be identical in both directions of every dyad. Where no symmetric unit exists, `rho` shall be reported as undefined and the relation assessed by paired ordinal comparison, explicitly demoted per FVP-2.

S49.6 — Both indices reported. Any report of network reciprocity shall carry `rho` (dyad-level) and `B` (node-level), together with the note that `1 − B >= rho`. Reporting only `B` is prohibited.

S49.7 — Numerator before denominator. No trust intervention shall be commissioned in a system whose measured `eta S − C` is negative. In that regime the intervention is predicted by Theorem 3.5 to worsen the equilibrium (§49.14). The order of operations is: raise `rho` and `eta`, cut `C`, confirm the numerator is positive, then invest in `a`.

S49.8 — Graduated sanction as default. Sanctions shall be graduated, with expulsion reserved and its price stated in advance as `ln N − ln(N−1)` of surplus ceiling plus the lost `Sym` volume. A proposal to expel shall carry that computation.

S49.9 — Nesting above the crossover. Where flat membership exceeds the computed crossover `n` at the measured `c_1`, the body shall be nested, and the nesting shall be a semilattice — units belonging to more than one parent — rather than a strict tree, per Alexander (1965). The tree figures in §49.13 are the cost floor, not the design.

S49.10 — Exit preserved. No obligation instrument shall bind a party who has not consented and may not withdraw. **Under CANON §5.4 this is not negotiable, and under §49.11 it is also the test that distinguishes a trust network from a capture network.** A network whose measured `a` is high and whose measured `rho` is near zero shall be reported as a capture structure.

S49.11 — No covert reliability testing. Sampling of verification shall be disclosed. Which commitments are sampled need not be.

S49.12 — Ledger privacy. Individual obligation balances shall not be published within the group. Distributions and dyad-level findings to the dyad only.

The Failure Modes

FM-49.1 — The Trust Programme in a Negative Regime. *Signature:* an engagement or culture initiative running in an organisation whose meetings produce no artefacts and whose coordination cost is unmeasured. Survey trust scores rise; nothing else does. *Damages:* wastes the budget and lowers `delta` while `eta S < C`. *Measurement:* compute `eta S − C` before commissioning. Protocol B for `rho`, Chapter 4's `eta = dQ_attributable / S_input` for `eta`. *Fix:* S49.7. Numerator first. *Confusable with:* FM-49.4.

FM-49.2 — Scalar Trust. *Signature:* a dashboard with one trust number, usually from an annual survey, usually on a five-point ordinal scale, usually averaged. *Damages:* destroys the domain structure that makes a actionable; produces uniform policy where differentiated policy is required. *Measurement:* attempt to reconstruct `a_ij^(d)` for three domains. If the pooled figure lies between them and predicts none of them, the diagnosis is confirmed. *Fix:* S49.1. Also note that averaging ordinal responses is prohibited by FVP-2 in the first place.

FM-49.3 — Verification Creep. *Signature:* a breach occurred; controls were added; the controls were never removed; `C` has been rising for three years and nobody can say by how much. *Damages:* `C`, permanently. Frequently pushes `h` negative on its own. *Measurement:* price the current `v` against $a*$ (S49.2). In most organisations that have never done this, `v` is far above the optimum for their measured `a`. *Fix:* recompute $a*$; sunset every control that was added in response to a specific incident more than `1/delta` periods ago and has not been re-justified. *Confusable with:* legitimate monitoring under Ostrom P4, which lowers `c_v`. The distinguishing test is whether the monitors are the users. If monitoring is done by a separate function, it is a `C` line; if it is a by-product of use, it is a `c_v` reduction.

FM-49.4 — Reciprocity Theatre. *Signature:* high interaction volume, celebrated. Newsletters about collaboration. `rho` unmeasured. *Damages:* `eta`, invisibly. The dashboard measures `Sym` and the store depends on `rho`. *Measurement:* Protocol B. In the worked consortium, volume of 68 units per quarter concealed `rho = 0.38`. *Fix:* publish the matched fraction alongside the volume, always.

FM-49.5 — The Node-Balance Illusion. *Signature:* a network reported as reciprocal on the strength of everyone's give and take roughly matching in aggregate. *Damages:* conceals dyad-level one-directionality; `1 − B` overstates `rho`. *Measurement:* compute both. In the worked case, 0.4412 against 0.3824. *Fix:* S49.6.

FM-49.6 — Window Shopping. *Signature:* a reciprocity figure that improved after the analyst "refined the measurement period." *Damages:* the credibility of the whole measurement. *Measurement:* ask when the window was chosen relative to seeing the data. *Fix:* S49.4, and report the sensitivity curve.

FM-49.7 — The Immortal Commitment. *Signature:* a rule everyone follows that nobody can date, justify, or find the author of. Usually discovered when a newcomer asks why. *Damages:* `C`, and `u` — because an unexamined rule is a framing nobody is permitted to contradict. *Measurement:* sample the commitment population and score each for last review date. A distribution with a long right tail is the signature. *Fix:* S49.3. This is `delta` engineering in the direction Chapter 3 warns is sometimes desirable: **a system with no forgetting cannot learn.**

FM-49.8 — Capture Mistaken for Culture. *Signature:* very high compliance, very high measured reliability, and members who do not leave. Flows are one-directional toward a small number of nodes. Departure is described as betrayal. *Damages:* everything, and the model will not show it if you only measure `a`. *Measurement:* `rho` and the sign structure of `b_i` (Protocol C, step 4). High `a`, low `rho`, and a small set of large net creditors is the signature. *Fix:* there is no parameter fix. **This is the failure mode the framework refuses to optimise** (CANON §5.4). The finding is a finding, and the required action is to restore exit.

FM-49.9 — Expulsion as First Resort. *Signature:* a governance body whose only sanction is removal. *Damages:* `N`, at `ln N − ln(N−1)` per removal, charged whether or not the removal was justified. *Measurement:* the table in §49.12: -7.7 per cent of `eta S` for one of eight. *Fix:* S49.8. Build the intermediate rungs before you need them; a sanction ladder cannot be constructed during the incident it is needed for.

FM-49.10 — Flat Above Crossover. *Signature:* a body of sixty-plus that still meets as one room and still decides as one room, and whose members describe the meetings as necessary and useless. *Damages:* `C`, quadratically. *Measurement:* estimate `c_1` from the observed per-path cost and find the crossover (§49.13). *Fix:* S49.9, with the semilattice caveat.

Worked Example — An Eight-Firm Regional Supplier Consortium

A consortium of eight independent manufacturing firms in one region, formed to share tooling, qualified staff, emergency capacity and test equipment. It has existed nine quarters. Membership is voluntary and members may leave with one quarter's notice. **Two members, firms 1 and 2, are substantially larger than the rest.**

The complaint brought to the engineer: “Participation is good but the consortium never seems to get anywhere. Every quarter we agree things and every quarter we start over.”

1. Encoding

- Period: one quarter.
- `N = 8` firms. Measured non-parallelism across their capability and market framings: `u = 0.52`. Hence `S = ln(1 + 7 * 0.52) = 1.5347`, against a ceiling of `ln 8 = 2.0794`. The consortium is operating at 73.8 per cent of its surplus ceiling, which is respectable and *not* the fault.
- `eta_max = 0.45` — the conversion the consortium would achieve at perfect balance, estimated from the three shared assets that did become permanent in the first two quarters.
- `C_0 = 0.30` per quarter: secretariat, meeting time, legal.
- `delta_0 = 0.05` per quarter: genuine irreducible loss, mostly staff movement inside member firms.

2. Baseline measurement

Protocol B and C, one quarter’s flow matrix, in units of documented favours (loaned tooling days, seconded staff days, test slots, referrals), all symmetric and ratio-scaled. Rows are givers:

to:	F1	F2	F3	F4	F5	F6	F7	F8
F1	–	1	1	0	0	0	0	0
F2	1	–	1	1	0	0	0	0
F3	6	7	–	3	2	1	1	0
F4	4	5	2	–	2	1	0	0
F5	3	4	1	1	–	1	0	0
F6	2	3	1	1	1	–	1	0
F7	2	2	0	0	0	1	–	1
F8	1	2	0	0	0	0	1	–

Totals: directed volume 68 favour-units per quarter; matched volume 26.

`rho_group = 26 / 68 = 0.3824`

Node balances `b_i` (given minus received):

firm	F1	F2	F3	F4	F5	F6	F7	F8
b_i	**-17**	**-21**	+14	+8	+5	+5	+3	+3

`B = 76 / (2 * 68) = 0.5588`, so `1 − B = 0.4412`. As §49.10 requires, note that `1 − B = 0.4412` exceeds `rho = 0.3824`: the node index flatters the network by fifteen per cent.

The sign structure is unambiguous. Two large net receivers, six net givers, and the worst dyads are one-directional:

dyad	x	y	volume	rho
(1,4)	0	4	4	0.0000

dyad	x	y	volume	rho
(1,5)	0	3	3	0.0000
(2,5)	0	4	4	0.0000
(2,6)	0	3	3	0.0000
(2,3)	1	7	8	0.2500
(1,3)	1	6	7	0.2857

This is Putnam's vertical patron-client pattern, measured. It is not what anyone intended; firms 1 and 2 are the largest and were the natural first call, and nine quarters of natural first calls produced a star.

Protocol A, reliance rate. Domain: commitments to supply capacity when asked. Sixty-two low-consequence commitments over four quarters, of which 21 were left unverified by the sampling rule. Honoured unverified: 13. `a_hat = 13/21 = 0.619`, taken as **`a = 0.62`**.

3. The diagnosis

```
eta   = eta_max * rho   = 0.45 * 0.3824 = 0.1721
delta = delta_0 + (1-a) = 0.05 + 0.38   = 0.4300
S     = 1.5347          eta*S = 0.2641          C = 0.30

Q_inf = (0.2641 - 0.30) / 0.4300 = -0.0836
```

The steady state is negative. By Theorem 3.5 this is the theorem stating its own failure mode: the regime is unsustainable, or the linear model must be replaced. The complaint — "we agree things and every quarter we start over" — is `Q_inf < 0` in plain English.

Headroom at the current `Q = 1.20`:

`h = eta ln N - delta Q - C = 0.3578 - 0.5160 - 0.3000 = -0.4582`

Negative. And by Corollary 3.4, **no arrangement of the interaction geometry rescues it.** The consortium could raise u to 1.0 and still fail, because the ceiling `eta ln N = 0.3578` is below `delta Q + C = 0.8160` on its own.

Trajectory from `Q_0 = 1.20`, no intervention:

quarter	0	1	2	3	4	5	6
Q	1.2000	0.6481	0.3335	0.1542	0.0520	-0.0063	-0.0395

Admissibility (`C_t <= (1-delta)Q_t + eta S_t`) requires `Q >= (0.30 - 0.2641)/0.57 = 0.0630`. **The trajectory breaches that between quarter 3 and quarter 4.** What the model predicts past that point is not negative stored order — that is meaningless — but the exit of the model's validity, which in this system means dissolution. The consortium had, on this measurement, about four quarters left.

4. What is required

By S49.7, the numerator comes first. The target: get `eta S − C` comfortably positive, then address a.

Required `eta` for a positive numerator at current S and C: `eta > 0.30/1.5347 = 0.1955`, so `rho > 0.4344`. That is the survival threshold and it is *barely* above the current 0.3824. **The consortium is 5.2 points of reciprocity away from not dying, and 51 points away from working.**

5. The intervention

Three changes, each attached to a parameter.

I-1 (raises `rho`): the reciprocal-call rule. Firms 1 and 2 commit that for each favour received they will originate one outbound offer to the same partner within the quarter, drawn from a published list of what they can supply. This is Ostrom P2 — congruence, benefits proportional to costs — implemented literally as a balance condition on the flow matrix.

I-2 (raises `rho` and volume): the capability register. Every firm publishes what it can lend. The small firms did not know what the large ones had. This raises `Sym` on the thin dyads and, being a published artefact, is also an `eta` device in Chapter 4's sense: it is structure that persists.

I-3 (raises a): the standing-commitment ledger with quarterly sunset. Every commitment recorded with a date and a twelve-month review. Per S49.3 this is deliberately a `delta`-raising device on obsolete items *and* an a-raising device on live ones, because most of the unhonoured unverified commitments in the baseline turned out to be commitments neither party could state precisely.

Deliberately **not** done: expulsion of the two large firms, or any sanction. Per §49.12 that would cost 7.7 per cent of `eta S` and, per FM-49.9, the ladder should exist before it is needed.

6. Post-intervention measurement, quarter 4 after

New flow matrix, same units:

to:	F1	F2	F3	F4	F5	F6	F7	F8
F1	–	3	4	3	2	1	1	1
F2	3	–	5	4	3	2	2	1
F3	6	7	–	3	2	1	1	1
F4	4	5	3	–	2	1	1	1
F5	3	4	2	2	–	1	1	1
F6	2	3	1	1	1	–	1	1
F7	2	2	1	1	1	1	–	1
F8	1	2	1	1	1	1	1	–

Directed volume 114 (up from 68); matched volume 102.

`rho_group = 102 / 114 = 0.8947` `B = 0.1053, 1 − B = 0.8947`

Balances are now `[−6, −6, +4, +2, +2, +2, +1, +1]`. Reliance rate re-measured by Protocol A after three quarters of the ledger: `a = 0.88`.

`eta = 0.45 * 0.8947 = 0.4026`

```
delta = 0.05 + 0.12   = 0.1700
eta*S = 0.4026 * 1.5347 = 0.6179     C = 0.30

Q_inf = (0.6179 - 0.30)/0.1700 = +1.8701
h at Q=1.20 = 0.8372 - 0.2040 - 0.3000 = +0.3332
half-life = 3.72 quarters (was 1.23); tau = 5.88 quarters (was 2.33)
```

Trajectory, both regimes, from the same `Q_0 = 1.20`:

quarter	0	2	4	6	8	10	12
before-regime Q	1.2000	0.3335	0.0520	-0.0395	-0.0693	-0.0789	-0.0821
after-regime Q	1.2000	1.4085	1.5521	1.6510	1.7192	1.7662	1.7985

7. Attribution — which change did the work

Per §49.14, decompose:

move	Q_inf	gain
baseline	-0.0836	—
reciprocity alone (I-1, I-2)	+0.7394	+0.8229
trust alone (I-3)	-0.2114	**-0.1278**
both	+1.8701	+1.9537
interaction		+1.2586

Had the consortium done I-3 alone — the ledger, the discipline, the reliability programme, which is the intervention every one of its members would have nominated — it would have measured better and performed worse. The reason is arithmetic and it is stated in Theorem 3.5: at `rho = 0.38` the numerator was negative, and reducing `delta` divides a negative number by less.

8. Falsification

Declared in advance, per Chapter 25 and FVP-6:

1. If `rho_group` does not exceed 0.4344 within three quarters, the intervention has failed to clear even the survival threshold and must be reported as unsuccessful.
2. If `rho` rises but the number of shared assets that survive a full year does not, then Postulate R — that matched volume is what converts — is wrong in this system and must be withdrawn, not defended.
3. If `a` rises while `rho` stays flat and `Q` nonetheless grows, the trust-as-delta mapping (Postulate T) is doing work the model says it cannot, and the model is wrong here.

4. **Shuffled-input control, mandatory.** Re-run the fit with the quarterly `U_t` sequence randomly permuted. If permuted `U` predicts `Q` as well as the real sequence, the entire causal story is a trend and must be reported as one.
5. If the two large firms' balances return to the baseline pattern within a year of the rule being relaxed, the improvement was in the rule and not in the relation, and the correct report is that `rho` was rented rather than built.

9. The rule, extracted

The consortium was not short of difference: at `u = 0.52` it was collecting 73.8 per cent of its surplus ceiling. It was not short of activity: sixty-eight favours a quarter among eight firms is a busy network. It was failing because two-thirds of that activity was one-directional, which put `eta` at 0.17, which put the numerator below the cost line, which made the steady state negative while every visible indicator — attendance, volume, goodwill — looked fine.

Trust is what you stop paying for. Reciprocity is what you get to keep. Obligation is the difference between the two directions, and it is the only one of the three that has a sign.

Measure all three or you will treat a conversion failure as a culture problem, spend a year on the culture problem, and arrive at a more reliable version of the same decline.

Chapter 50 . The Internal/External Correspondence

This is the central argument of Volume Two, and it is short enough to state in one line before the chapter defends it:

Psychology and sociology are not two subjects. They are one set of relations measured at two scopes.

Volume One took the storage law outward — rooms, tables, teams, institutions, cities. Volume Two turned the same instrument inward — attention, framing, character, capacity. Both sides of that turn have now been written with the same six parameters, the same ceiling, the same leaky integrator and the same failure taxonomy. This chapter takes the master table of CANON Addendum II §5 and develops it row by row, giving each parameter its internal measurement and its external measurement side by side, so that the correspondence is not asserted but exhibited.

And then it does the thing that a correspondence chapter usually does not do. **Section 50.14 is the longest section in this chapter, and it is the section that says where the correspondence fails.** A mind is not a small society and a society is not a large mind. There are at least nine specific places where the mapping breaks, most of them structural rather than incidental, and three of them severe enough that any result transferred across the scope boundary in their vicinity should be treated as unsupported until measured separately at the new scope.

A correspondence stated without its limits is not an argument. It is a slogan. The limits are in §50.14 and they are stated at least as forcefully as the correspondence itself.

The Law

50.0.1 What is proven

Exactly one thing, and it is a statement about form:

`Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t` `S_t = F(U_t) = ln(1 + (N−1)U_t)`

The equation contains no term that names a substrate. There is no variable for “person” and no variable for “organisation.” Every theorem proved about it — the exact solution (Thm 3.1), the one-step growth condition (Cor 3.2), the uniform bound (Thm 3.3), the sustainability ceiling (Cor 3.4), the steady state (Thm 3.5), the headroom gauge (§6.4) — is a theorem about a first-order linear recursion with a loss term, and it holds wherever that recursion holds.

CANON §1.4 states the consequence precisely and this chapter does not exceed it:

> The storage law is scale-free in form: only (delta, eta, N, C) change between a desk, a room, a team, a building, an institution, a city. Same equation, different parameters. **That is a claim about MODEL FORM, not about physics.**

Two further proven facts do the structural work in what follows:

- **Banach’s fixed point theorem.** For `delta > 0` the map is a contraction with modulus `(1 − delta)`, so a unique fixed point exists and iteration converges to it geometrically from any starting value. That is why the trajectories at both scopes have the same shape (§50.13).
- **Strict concavity of `f`** (Paper A, Thm 3.3) and the hard ceiling `0 <= F <= ln N` (Paper D, Cor 2.2). Both are properties of the logarithm and neither knows what `N` is counting.

50.0.2 What is postulated, badged

Thesis, badged as the author's thesis and not as a theorem: the internal and external scopes are the same relational system, and the parameters of one have measurable counterparts in the other.

That is stronger than the proven statement and it must be kept separate from it. The proven statement says the *equation* transfers. The thesis says the *phenomena* correspond — that forgetting and turnover are the same parameter at two scopes, that character and institution are the same store, that cognitive load and coordination overhead are the same dissipation.

A modelling thesis of this kind is not established by exhibiting a table. It is established by FVP-7: the same operationalisation, applied at both scopes, must recover parameters that stand in the predicted relation, and the prediction must be able to fail. §50.15 states exactly what would kill it. Until that programme is run, the correspondence is a research hypothesis with a good structural argument behind it and no cross-scope replication behind it, and this chapter says so in those words.

50.0.3 Honesty clauses

- As deployed, the surplus functional is a **lower bound** on the ensemble's information. Every quantity below is a floor.
- The triangle inequality **fails** for geometric triples (Paper A §5). Subadditivity holds on u only. **This matters here specifically: it means you cannot compose scopes metrically** — the distance from a mind to a team to an institution does not obey a triangle inequality, so multi-scope chains are not a metric hierarchy.
- Paper A is pure mathematics and makes **no empirical or metaphysical claim**. Surplus is not identified with Shannon entropy or mutual information (Paper D). **No thermodynamic negentropy law is proved** (Paper E).
- **Structural similarity is not identity.** This clause has already been used in Chapter 34 for the predictive-processing correspondence and it applies here with more force, because the temptation is greater: two systems described by the same equation are not thereby the same kind of system. Population dynamics and radioactive decay share an equation and share nothing else.

The Engineering

50.1 How to read the master table

The table below is CANON Addendum II §5, reproduced exactly. Each row gets a section. Each section states: what the parameter *is* at that scope, **how you measure it there**, the unit, and the one place where the two measurements diverge in kind rather than in value.

Term	INTERNAL (psychology, personal)	EXTERNAL (sociology, interpersonal)
N	distinguishable categories in one mind	distinct domains held in a group

Term	INTERNAL (psychology, personal)	EXTERNAL (sociology, interpersonal)
u	non-parallelism between one's OWN framings	non-parallelism between members
N_eff	how many genuinely different views you hold	how many genuinely different members
F	internal surplus: many windows, one tower	group surplus
Q	character, skill, memory, internalized structure	institution, culture, infrastructure
delta	forgetting, skill atrophy, drift of self	turnover, decay of practice, lost tradition
eta	how experience becomes character	how interaction becomes institution
C	cognitive load, effort, maintenance of self	coordination overhead, friction
h	personal capacity for the new	organizational headroom
M	the moment as lived	the situation as the group constructs it
D/P	frame-to-intake ratio	doctrine-to-data ratio

One reading instruction before the rows. **The correspondence is not that the two columns describe the same thing. It is that they occupy the same position in the same equation.** `delta` internal and `delta` external are different phenomena — forgetting is not turnover — but they enter the recursion identically, they are estimated by the same procedure, and they have the same consequences for the trajectory. That is the whole claim, and it is both weaker and more useful than the claim that a company is a mind.

50.2 Row `N` — the category count

Internal. `N` is the number of distinguishable categories the mind holds in the domain under analysis. Not the number of thoughts; the number of *kinds* it can tell apart. Chapter 39 developed this as the count over one observer's own framings.

Measurement, internal: a free-sort task. Present the person with 40–60 items from the domain and ask them to sort into as many groups as they find natural, with a name per group. `N` is the count of groups that survive a re-sort a week later. **The re-sort is the whole protocol**: categories that do not reappear were produced by the task, not held by the person. Report test-retest agreement alongside `N`; below about 0.6 the count is not a measurement.

External. `N` is the number of distinct domains the group holds — not headcount. Chapter 7 established this and Chapter 11 established how badly headcount overstates it.

Measurement, external: the capability inventory. Enumerate the domains in which at least one member can act without consulting anyone outside the group, verified by an independent instance of the member acting in that domain. `N` is the count of such domains.

Where the two diverge in kind. External `N` is enumerable from public evidence and can be

audited by a third party. Internal N is a modelling choice about how finely to carve, and two competent analysts will return different counts from the same person. **The external count has a fact of the matter that the internal count does not**, and this asymmetry runs the opposite way to most people's intuition. It is easier to count what a team can do than what a mind can distinguish.

50.3 Row u — non-parallelism

Internal. u is the degree to which the person's own framings of a situation are genuinely different rather than the same framing in different words. u = 1 − (x · y)^2 for a pair of framings; the group's value is the mean over pairs.

Measurement, internal: the independent-prediction test. Ask the person to produce k framings of one situation, then — separately, without re-reading — ask each framing for a prediction about a specific future observable. u is estimated from the disagreement structure of the predictions. **Framings that predict identically are the same framing.** This is the only defensible operationalisation, because framings stated in words are trivially made to look different and non-trivially made to be different.

External. Same instrument, different subjects: u is the mean pairwise non-parallelism between members' framings, measured by having them predict independently before discussing (Chapter 22, the pre-discussion round).

Where the two diverge in kind. Internally, all framings are held by one processor and share its priors; they are correlated by construction and u is bounded well below 1 in practice. Externally, u near 1 is achievable because members can be genuinely independently sourced. **The internal ceiling on u is lower than the external one, and Chapter 39's result — that the second genuinely different framing is the largest gain you will ever buy — is a consequence of that lower ceiling, not a restatement of the concavity.**

50.4 Row N_eff — the effective count

N_eff = 1 + (N−1)u at both scopes, without modification. This row is the cleanest correspondence in the table because it is pure arithmetic on the two rows above it.

With N = 7 at either scope:

u	N_eff	S = ln N_eff	share of ceiling ln 7 = 1.9459
0.15	1.900	0.6419	33.0%
0.35	3.100	1.1314	58.1%
0.55	4.300	1.4586	75.0%
0.75	5.500	1.7047	87.6%
0.95	6.700	1.9021	97.7%

Seven framings held by one person at u = 0.15 and seven people in a room at u = 0.15 are both worth 1.9 of anything. The table does not care which. That is the correspondence in its most literal and least contestable form: **the audit for redundancy is scope-independent**, because redundancy is a geometric property and geometry does not know what it is measuring.

Measurement, both scopes: compute from `N` and `u`. Never report `N_eff` from a separate instrument; it is a derived quantity and reporting it independently invites the double-counting that Chapter 11 catalogues.

50.5 Row `F` — surplus

Internal. Book One, Node 03 HORIZON: *the many windows of the one tower*. Chapter 39 derived it as Paper A applied within a single observer, and it is a derivation rather than an analogy — the theorem does not require the vectors to belong to different people.

External. The group's surplus per interaction, Chapter 7.

Ceiling at both scopes: `ln N`. **And here the correspondence produces a result that neither scope would produce alone.** CANON Addendum II §1 records the forced consequence: since the Moment normalises into `[0, ln N]` to serve as `S_t`, the maximum possible moment is `ln N`, which is Ashby's Law of Requisite Variety (1956) appearing in first-person experience. **The richest available moment is bounded by the log of how many kinds of thing the perceiver can distinguish.** That statement was not chosen. It follows from putting the internal and external ceilings in the same equation, and it is the strongest single piece of evidence that the two scopes are running the same machinery.

Measurement, both scopes: `F = ln(1 + (N−1)u)`, derived, never measured directly.

50.6 Row `Q` — stored order

Internal. Character, skill, memory, internalised structure. What persists in the person when the situation that built it has gone.

Measurement, internal: the unprompted-retention probe. Sample the domain, wait one period without input, and test performance without cues. `Q` is scored on the retained-capability scale defined in Chapter 40. **The essential feature is the absence of prompting**: capability that appears only when scaffolded is not in `Q`, it is in the scaffold, which belongs to the environment.

External. Institution, culture, infrastructure. What persists in the group when the people who built it have gone.

Measurement, external: the identical probe with the identical essential feature — the substitution test. Replace the individual and see whether the capability survives. Star & Ruhleder (1996) supply the diagnostic principle for both scopes: **infrastructure becomes visible only on breakdown**, so `Q` is most accurately measured during a disruption and is systematically underestimated during normal operation.

Where the two diverge in kind. Internal `Q` has a hard terminus. External `Q` does not: an institution's store can in principle outlive every person who has ever held it, and often does. **A mind's store is bounded by a lifetime and a group's is not, which means the two scopes have different upper limits on the time constant `1/delta` and therefore different meaningful ranges for `delta` itself.** This is the first crack, and §50.14 develops it.

50.7 Row `delta` — retention loss

Internal. Forgetting, skill atrophy, drift of the self away from what it had settled into.

Measurement, internal: Chapter 3's estimator, unchanged. Measure `Q` at two times under zero input; `delta = 1 − Q_1/Q_0`. Half-life `t_half = ln 2 / −ln(1 − delta)`.

External. Turnover, decay of practice, lost tradition. And here the external scope has a term with no internal counterpart, which must be composed rather than assumed:

```
delta_total = 1 − (1 − delta_internal)(1 − tau)
```

where `tau` is the per-period turnover rate. The two losses are independent, so their survivals multiply.

`delta_int`	`tau`	`delta_total`	`Q_inf` relative to `tau = 0`
0.05	0.00	0.0500	100.0%
0.05	0.05	0.0975	51.3%
0.05	0.10	0.1450	34.5%
0.05	0.20	0.2400	20.8%
0.05	0.30	0.3350	14.9%
0.05	0.50	0.5250	9.5%

Ten per cent turnover cuts steady-state stored order to a third, and it does so before anyone has written a single bad document. This row is the one where the external scope is strictly harder than the internal one: your knowledge does not resign.

The measurement caution both scopes share: `delta` is per period, and the period lengths at the two scopes typically differ by one to two orders of magnitude. Any cross-scope comparison of `delta` values without normalisation is meaningless. §50.14, breakdown B7, computes exactly how meaningless.

50.8 Row `eta` — conversion

Internal. How experience becomes character. Chapter 40's join: the Moment equation produces the input and `eta` decides how much of it stays.

Measurement, internal: `eta = dQ_attributable / M_input`. Requires the retention probe of §50.6 and an independent estimate of the input. It is the hardest parameter to measure at either scope, because the attribution is the whole difficulty.

External. How interaction becomes institution. Chapter 4.

Measurement, external: `eta = dQ_attributable / S_input`. Same estimator, same difficulty, one advantage: the artefacts are countable. A decision log exists or does not. **Meetings with no artefact have `eta` approximately 0**, and that is checkable from outside in a way that "did that experience change you" is not.

Where the two diverge in kind — and this is a real asymmetry, not a measurement inconvenience. Chapter 49 established that at group scope `eta = eta_max * rho`: only the reciprocated fraction of interaction converts, because structure only one party holds is a debt rather than a structure. **There is no internal analogue of `rho`, because your framings cannot renege.** A framing does not have interests, cannot defect, cannot fail to return a favour. The entire apparatus of reciprocity, sanction and obligation — Chapter 49 in its entirety — transfers inward only as metaphor. §50.14, breakdown B4.

50.9 Row C — dissipation

Internal. Cognitive load, effort, the ongoing maintenance of a self: the cost of holding the structure up, paid every period whether or not anything is built.

Measurement, internal: dual-task interference, or time-on-maintenance sampling — the fraction of the period spent keeping the existing store in order rather than adding to it. Both are indirect. The honest position is that internal C is currently the least well-measured parameter in the table.

External. Coordination overhead, friction, conflict, compliance, rent, commute. Chapter 5.

Measurement, external: direct. Coordination cost is a line item. And it has a known scaling law that the internal scope does not obviously share — Brooks (1975): communication paths grow as n(n−1)/2, so **cost grows quadratically while surplus grows logarithmically**, giving a computable crossover beyond which no arrangement helps.

Where the two diverge in kind. External C has that quadratic term because members must coordinate pairwise. Internally the framings do not have to negotiate with each other in the same way — a mind is not paying n(n−1)/2 in inter-framing meetings. **The external scope has a cost law with an exponent; the internal scope, as far as the framework can currently say, does not, and this is a genuine structural difference rather than a difference in parameter value.**

50.10 Row h — headroom

h = eta ln N − delta Q − C at both scopes, unmodified. The three regimes are identical:

```
h < 0   unsustainable in this model. No geometry saves it.  (Cor 3.4)
h = 0   maintenance only. Hold, do not build.
h > 0  growth available; any requested increment r must satisfy r <= h. (Cor 4.2)
```

Internal: personal capacity for the new. The reason a person at high Q and high C cannot take on one more thing is not weakness of will; it is h < 0, and Chapter 6's structural result holds identically at both scopes: **success consumes its own headroom**, because delta Q grows with Q.

External: organisational headroom. The ossification of successful institutions is the same sentence with different nouns.

Measurement, both scopes: derived from the four parameters. Never surveyed. "Do you have capacity" is a question about perceived h, and perceived h is exactly the quantity that Chapter 48's doctrine-to-data fault distorts.

This row is where the correspondence pays for itself most directly, because **the same diagnostic question — is the maintenance term eating the ceiling — is asked identically of a person and of a department, and answered by the same arithmetic.**

50.11 Row M — the moment

Internal. M = (D/P) . E, the moment as lived. Chapter 34.

External. The situation as the group constructs it. Chapter 47 built this at group scope with D = shared doctrine, P = collective perception, E = environment.

Measurement, both scopes: the four limiting cases are the diagnostic instrument, and they are identical in structure:

```
P -> 0    M -> infinity   frame with no intake   internal: dogma, apophenia
                                                 external: an institution that has stopped listening
D -> 0   M -> 0       intake with no frame   internal: sensation without experience
                                             external: a body that measures everything and
                                                                concludes nothing
E -> 0    M -> 0          no environment         both: no situation, no moment
D = P     M = E             THE FIXED POINT          both: undistorted contact
```

The fixed point `D = P` is the same fixed point at both scopes, and it means the same thing: when the frame matches the intake, the moment equals the environment. That is veridical perception for a person and accurate situational assessment for an organisation, and it is one equation.

50.12 Row `D/P` — the calibration ratio

`r = D/P`, distortion `Delta = |ln r|`, in the same logarithmic register as everything else in the framework.

Internal: frame-to-intake. `r > 1` over-framed (projection, seeing the expected); `r < 1` under-framed (flooding, cannot make it mean). Chapter 35.

External: doctrine-to-data. **An over-framed institution is one whose doctrine outruns its intelligence — it sees what its model predicts. An under-framed institution drowns in data it cannot read.** Chapter 48. Both are measurable and both are common.

Bands, at both scopes:

`|ln r|`	`r` (over)	`r` (under)	reading
< 0.10	< 1.105	> 0.905	calibrated
0.10–0.40	1.105–1.492	0.905–0.670	drift
0.40–1.00	1.492–2.718	0.670–0.368	fault
> 1.00	> 2.718	< 0.368	severe

Measurement, both scopes: the prediction-surprise ledger. Record predictions with confidence before the observation, score them after, and estimate `r` from the ratio of confidence held to information actually supplied. Over-framing shows as high confidence with flat accuracy; under-framing shows as low confidence with accuracy that would have supported more.

Where the two diverge in kind. An institution's `D` is a public document and its `P` is a measurable throughput. A person's `D` and `P` are both private. **The external version of this row is the more measurable one, which is the reverse of the intuitive expectation and is worth stating plainly: doctrine-to-data is easier to measure than frame-to-intake, because doctrine is written down.**

Two further things transfer with this row, and both correct the way `r = 1` reads in the band table above.

First: at the fixed point, headroom is exactly zero — at either scope. Put `r = 1` into the contact relation and the input sits at its ceiling, `S = ln N`. Paper E Theorem 3.5 then gives `Q_inf = (eta.ln N − C)/delta`, and the master gauge of row h reads

```
h = eta ln N − delta Q_inf − C = eta ln N − (eta ln N − C) − C = 0
```

with `delta` and `C` cancelling out. **Perfect contact drives the input to its maximum, the maximum input drives the store to its maximum, and the maximum store spends the entire budget on maintenance.** Chapter 34 §34.16 derives this for a person and Chapter 47 §47.6 for an institution, and that it comes out identically at both is among the cleaner pieces of evidence for this chapter's thesis: it is a cancellation in the algebra, not a resemblance in the prose. **The consequence, at both scopes: `r = 1` is not an unqualified target.** It maximises fidelity and veridical throughput, and it leaves nothing to build with. Room at a settled state requires shedding `Q` deliberately (row `Q`; Chapter 23 externally, Chapter 40 §40.9 internally) or holding a declared `r != 1` as policy.

At external scope the fixed point acquires an exact join to Paper A, derived in Chapter 47 §47.8:

```
ln r = ln D − ln p_bar − F
```

The group's surplus enters the doctrine-to-data ratio one nat for one nat. Calibration becomes a geometry requirement rather than an attitude: a body carrying doctrine load `D/p_bar` cannot reach `r = 1` at all unless `D/p_bar <= N`, and the residual floor `Δ_min = ln(D/p_bar) − ln N` is where heavy institutions live. **This is the exact seam between Volume Two's two governing equations, and it is available only at external scope** — the internal analogue would need a decomposition of one person's intake into countable channels, and §50.14's aggregation disanalogy says that is not on offer. Record it as a place where the correspondence is **asymmetric in what can be computed**, not merely in what can be measured.

Second: the ledger this row drives must be kept in two tracks, at both scopes. Chapter 34 §34.21 establishes that over-framing **increases** stored order — the storage law integrates the moment as lived and contains no term asking whether it was earned — and Chapter 40 §40.2 turns that into an instrument: `Q` against veridical `Q^V`, with stored error `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` (the cost term cancels exactly) and correction load `C_t = c_0 + kappa.Z_t`. **Dogma banks faster than knowledge, and the bill arrives on a lag.** Because `V = E.min(r,1)` is flat above the fixed point while `M = rE` keeps climbing, **anyone optimising raw `Q` is driven away from `r = 1`, actively and by arithmetic**: a person accumulating a confident self-account, an institution accumulating doctrinal coverage, identically. The knife edge is `kappa = delta` — below it, raw `Q_inf` rises without bound in `r`. **Row `Q` of the master table should therefore be read as two rows**, and a report of total store without `Z` beside it is a report of a number that grows fastest when the reporter is most wrong. Chapter 47 §47.9.1 works this at group scope; Chapter 48 gives the corrections at both.

50.13 What the eleven rows share: one trajectory shape

Every row above is a parameter of the same recursion, and the recursion is a contraction whenever `delta > 0`. Banach therefore guarantees that both scopes converge geometrically to their own fixed points, and that **after normalising time by the time constant `tau = 1/delta`, the two trajectories are the same curve.**

Computed. Internal: period one day, `delta = 0.02`, `eta = 0.30`, `M = 1.20`, `C = 0.15`, so `tau`

= 50 days and Q_inf = 10.5000. External: period one quarter, delta = 0.16, same eta, S and C, so tau = 6.25 quarters and Q_inf = 1.3125. Both from Q_0 = 0:

t / tau	internal Q/Q_inf	external Q/Q_inf	continuous limit 1 − e^−k
0.5	0.3965	0.4073	0.3935
1.0	0.6358	0.6487	0.6321
2.0	0.8674	0.8766	0.8647
3.0	0.9517	0.9636	0.9502
4.0	0.9824	0.9872	0.9817
5.0	0.9936	0.9955	0.9933

The two columns agree to within two per cent at every point, and both track the continuous exponential. **The steady states differ by a factor of eight and the clocks differ by a factor of about a hundred and forty, and the shape is identical.** That is the correspondence's strongest demonstrable content: **the form transfers, the clock does not, and the levels do not.**

Book One's node 10 FIXEDPOINT is Paper E Theorem 3.5 at internal scope, and Chapter 23's institutional equilibrium is Theorem 3.5 at external scope. **They are the same theorem, and the reason arrival feels like recognition rather than victory is that a contraction arrives where it was always going to arrive** — at either scope.

50.14 Where the correspondence breaks down

Nine failures: three severe, four structural but manageable, one a matter of measurement practice, and one absolute. Each is stated as sharply as it can be stated, with its computation where it has one, and with the transfer restriction it implies.

B1 — Exit. A group's members can leave; a mind's parts cannot. SEVERE. This is the deepest disanalogy and everything else in this section is downstream of it in some way.

A framing does not resign. A category does not withdraw its labour. A memory does not accept an offer from a competitor. But every member of a group is a separately deciding agent with an outside option, and that fact changes the *type* of the model, not just its parameters.

Consequence 1: N is endogenous externally and exogenous internally. In the internal model N is set by the analyst's carving and stays fixed while the trajectory runs. In the external model N is an outcome of the trajectory: a group whose Q falls loses members, which lowers N, which lowers the ceiling eta ln N, which lowers Q further. **The external system has a feedback loop through N that the internal system does not have at all.** Every theorem in Paper E is proved for fixed N. None of them is proved for the closed loop.

Consequence 2: group-level sustainability does not imply member-level sustainability. Headroom is computed for the body. Each member has their own:

```
h_i = eta ln N / N − delta Q / N − c_i
```

Computed for `N = 9`, `u = 0.6`, `eta = 0.35`, `delta = 0.08`, `Q = 4.0`, with heterogeneous cost shares. Per-member share of the ceiling is `0.08545`; per-member share of maintenance is `0.03556`; so the critical private cost is `c_i = 0.04989`:

member	c_i	h_i	status
1	0.02	+0.02989	ok
2	0.03	+0.01989	ok
3	0.05	-0.00011	unsustainable
4	0.09	-0.04011	unsustainable
5	0.14	-0.09011	unsustainable
6	0.22	-0.17011	unsustainable

Four of six members are individually below water in a body that may well be reporting positive aggregate headroom. Those four will leave, and they will leave in cost order — the highest-cost members first, which are frequently the members carrying the highest u, because carrying a divergent framing in a converged room is expensive. **The internal model has no equation that can even express this, because a mind's subsystem cannot have a private balance sheet.**

Transfer restriction: **no result derived at internal scope may be transferred outward without checking member-level headroom separately.** The aggregate is not sufficient.

Named in the literature: Hirschman, *Exit, Voice, and Loyalty* (1970), is the standing treatment of the three options available to a dissatisfied member of an organisation. Two of the three — exit and loyalty as a mediator of exit — have no internal counterpart at all.

B2 — Heterogeneity. A group's parameters are distributions; a scalar model is a mean field. SEVERE. The storage law as written has one `delta`, one `eta`, one `C`. A group has one per member. Writing a single value is a **mean-field approximation**, and mean-field approximations fail in a specific, computable way when the underlying distribution is not tight.

Computation 1: the aggregate `delta` is not any member's `delta`, and it is not even constant. Take a store split evenly between two compartments, `Q_A = 50` at `delta_A = 0.05` and `Q_B = 50` at `delta_B = 0.50`. The Q-weighted mean is `delta_bar = 0.275`. Run both, with no input:

t	true total	mean-field prediction	error	error %	implied aggregate `delta`
0	100.0000	100.0000	0.0000	0.00%	—
1	72.5000	72.5000	0.0000	0.00%	0.2750
2	57.6250	52.5625	+5.0625	+8.79%	0.2052
3	49.1187	38.1078	+11.0109	+22.42%	0.1476
5	40.2515	20.0304	+20.2211	+50.24%	0.0821
8	33.3663	7.6332	+25.7332	+77.12%	0.0550
12	27.0302	2.1089	+24.9213	+92.20%	0.0504

The mean-field model is exactly right for one period and wrong by 92 per cent after twelve. The reason is structural: the fast compartment empties, so the surviving store is increasingly the slow one, and the implied aggregate `delta` drifts from 0.275 down toward 0.05. **A single `delta` fitted to early data will catastrophically over-predict decay, and a single `delta` fitted to late data will catastrophically under-predict it.** Neither is a parameter of the system; both are artefacts of when you looked.

Computation 2: aggregate u does not identify the structure, and the structures differ in fragility. Model each member as `v_i = cos(th_i) e_0 + sin(th_i) f_i`, with `e_0` the shared direction and the `f_i` mutually orthonormal private directions. Then `u_ij = 1 − c_i c_j` with `c_i = cos^2(th_i)`, and the mean pairwise u follows in closed form.

configuration	composition	mean u	S at N = 9
A, homogeneous	nine members all at c = 0.6325	0.5999	1.7578
B, heterogeneous	six at c = 0.95, three at c = 0.0501	0.5999	1.7578

Identical aggregate u. Identical surplus. Now remove three members:

configuration	which three	remaining u	remaining S at N = 6
A	any three	0.5999	1.3862
B	the three divergent	0.0975	**0.3971**
B	three conformists	0.7904	1.5998

Configuration A is departure-invariant in u. **Configuration B loses 77 per cent of its surplus to the departure of three of nine, where configuration A loses 21 per cent** — and B *gains* if the right three leave instead. The two systems are indistinguishable from the aggregate and behave three and a half times differently under a shock.

This is a failure of FVP-5, identifiability, and it is a failure that only exists at external scope, because it is a failure about who leaves — which is B1 again.

Transfer restriction: **at external scope, u shall be reported with its distribution across members, never as a scalar alone.** At internal scope the scalar is defensible because the framings cannot separate.

B3 — Aggregation is not always well defined. SEVERE. The master table's `D/P` row assumes that a group *has* a doctrine — that there is a coherent `D` to put in a numerator. For a single mind this is usually approximately true. For a group it can be false as a matter of arithmetic, not of sloppiness.

Three members, three options, these preference orders:

```
member 1:  A > B > C
member 2:  B > C > A
member 3:  C > A > B
```

Pairwise majorities, computed:

```
A vs B:  2–1  ->  A
B vs C:  2–1  ->  B
A vs C:  1–2  ->  C
```

The group prefers A to B, B to C, and C to A. There is no transitive collective ordering to extract. This is the Condorcet paradox (1785), and Arrow's impossibility theorem (1951) generalises it: no rule aggregating three or more ordinal preferences can simultaneously satisfy unrestricted domain, non-dictatorship, the Pareto criterion, and independence of irrelevant alternatives.

The consequence for this chapter is exact: the group's `D` may not exist as a well-defined object, in which case `D/P` is not merely hard to measure but undefined, and Chapter 48's diagnostic does not apply. A group can be cyclically incoherent while every member is individually coherent.

The honest counter-weight, stated because the disanalogy should not be overclaimed: individual preferences also exhibit intransitivity — preference reversals under reframing are among the more robust findings in judgement research. So the disanalogy is one of degree and of mechanism rather than of kind. **But the mechanism differs in a way that matters: individual intransitivity is an error, and collective intransitivity can be the correct output of a faultless procedure applied to faultless inputs.** No amount of member-level rationality removes it.

Transfer restriction: **before computing a group `D/P`, test for cyclicity.** If the group's collective ordering over the relevant options is cyclic, report `D` as undefined and diagnose by member-level `D_i/P_i` distributions instead.

B4 — Reciprocity has no internal analogue. The homunculus trap. STRUCTURAL. §50.8 stated it: `eta = eta_max * rho` externally, and internally there is no `rho`, because a framing cannot renege. Extending the machinery of Chapter 49 inward requires positing internal agents with interests — a homunculus, and then a second homunculus to explain the first.

The trap is specific and it is easy to fall into precisely because the equations permit it. The storage law will accept an "internal reciprocity" parameter without complaint; the equation has no way to object. What objects is the measurement: `rho` requires a flow matrix `x_ij` between agents, and there are no internal agents whose flows can be counted.

The partial exception, stated honestly: intertemporal-choice research does model a person as a sequence of selves whose interests conflict, and the technology of precommitment — binding your later self, the Ulysses contract — has an obvious structural resemblance to a contract between parties. This is a real modelling tradition and it is not obviously wrong. **But the resemblance is to contract, not to reciprocity: the later self did not choose to enter the arrangement, and cannot decline it, which is B1 again.** A framework whose fourth refusal is *never remove a person's ability to decline* must notice that precommitment is precisely the removal of a future person's ability to decline, and that this is permitted only because the two persons are the same person.

Transfer restriction: **Chapter 49 transfers outward-to-inward as metaphor only.** Nothing in that chapter's arithmetic — `rho`, the obligation ledger, the sanction ladder, `a*` — has an internal measurement, and none of it should be given one.

B5 — The observation asymmetry. Different error structures at the two scopes. STRUCTURAL. Internal measurement has privileged access and unreliable instruments: introspection is available continuously, at no cost, and is **biased in a direction correlated with the quantity being measured**. A person who is over-framed (`r > 1`) is by definition seeing what their frame predicts, including when they look at their own frame. The error is not noise; it is signal-dependent distortion.

External measurement has public access and sparse, noisy instruments: you observe behaviour, at cost, at intervals, through selection filters. **The error is large but comparatively unbiased**, or at least biased in ways that are themselves observable.

The consequence is that the same parameter has different identifiability profiles at the two scopes (FVP-5), and therefore requires different estimators — not the same estimator with a different sample size. Internal `C` is the clearest case: §50.9 conceded that it is the least well-measured quantity in the table, and the reason is that the cost of maintaining a self is paid by the same apparatus that would report it.

Transfer restriction: **an operationalisation validated at one scope is not thereby validated at the other.** FVP-7 must be re-run per scope. This is not pedantry; it is the specific gate that CANON Addendum II §6 identifies as the one most often skipped.

B6 — Boundary definition runs the opposite way at the two scopes. STRUCTURAL. Noted at §50.2 and §50.12 and collected here because it is systematic and counter-intuitive:

quantity	easier to measure internally	easier to measure externally
`N`		yes — capability inventory is auditable
`u`		yes — independent prediction rounds
`Q`		yes — substitution test
`delta`	yes — no turnover term to disentangle	
`eta`		yes — artefacts are countable
`C`		yes — coordination cost is a line item
`M`	yes — you are the one having it	
`D/P`		yes — doctrine is written down

Six of eight are easier externally. The common assumption is the reverse — that we know our own minds and can only guess at institutions. **The framework's measurement requirements invert that assumption almost completely, and this is one of the more useful things the correspondence produces**: if you want to learn how to measure an internal parameter, look at how the external version is measured, because the external version is usually the better-instrumented one.

B7 — Timescale. `delta` is per period and the periods differ by two orders of magnitude. PRACTICE. The scope transform is

```
delta_T = 1 - (1 - delta)^T
```

for aggregating T short periods into one long one. It is exact and it is routinely omitted.

`delta` per period	`T = 7`	`T = 30`	`T = 52`	`T = 365`
0.001	0.0070	0.0296	0.0507	0.3059
0.005	0.0345	0.1396	0.2295	0.8395
0.010	0.0679	0.2603	0.4070	0.9745
0.020	0.1319	0.4545	0.6503	0.9994
0.050	0.3017	0.7854	0.9306	1.0000
0.100	0.5217	0.9576	0.9958	1.0000

And the trap, stated with its number. A `delta` of 0.02 at internal scope with a daily period, and a `delta` of 0.02 at external scope with a quarterly period, look like the same parameter value:

```
internal 0.02/day  compounded over 91 days   ->  0.8409 per quarter
external 0.02/quarter expressed per day      ->  0.000222 per day
ratio of implied daily rates                 ->  90.1x
```

Two numbers that print identically differ by a factor of ninety. Every table in this book that compares scopes must either normalise or state its period, and this chapter's §50.13 does the former explicitly.

Transfer restriction: **`delta` and `eta` shall never be quoted without their period. `Q`, `C` and `h` inherit the same requirement.** Only `N`, `u`, `N_eff` and `F` are period-free.

B8 — Rival versus non-rival stored order. STRUCTURAL, and it has no internal counterpart. Whether members exit depends on what they get, and what they get depends on whether `Q` is rival.

Rival case (the store is divided): each member's per-period return is `eta ln(N_eff)/N - c_i`. With `eta = 0.35`, `c_i = 0.020`, outside option `w = 0.030`:

N	S = ln N	per-capita	return	decision
3	1.0986	0.12817	+0.10817	stay
6	1.7918	0.10452	+0.08452	stay
12	2.4849	0.07248	+0.05248	stay
20	2.9957	0.05243	+0.03243	stay
30	3.4012	0.03968	+0.01968	**exit**

ln(x)/x is maximised at x = e = 2.71828. **Under equal division of a logarithmic surplus, the per-capita return is maximised at a group of about three, and every additional member reduces everyone's share.** Rival stored order pushes groups toward the smallest size that clears the outside option.

Non-rival case (the store is a public good — a standard, a piece of infrastructure, a body of practice): each member's return is eta ln N − c_i, which increases without bound in N:

N	3	5	8	12	20	30
return	+0.3645	+0.5433	+0.7078	+0.8497	+1.0285	+1.1704

The same group, the same surplus, the same members — and the sign of the size gradient flips entirely depending on whether the stored order is divided or shared. This is a fork in the external model with no internal counterpart whatever, because a mind's stored order is not shared with anyone and the question does not arise.

Transfer restriction: **any external application of the storage law must declare whether Q is rival or non-rival before predicting anything about group size.** Predictions about optimal N transferred from internal reasoning are unsupported.

B9 — The ethical asymmetry. ABSOLUTE. The mathematics is symmetric. The permissions are not.

You may restructure your own N, force your own u, raise your own delta on a habit you want gone, and spend your own h however you like. **You may not do any of those things to another person, and the framework's fourth refusal is exactly this: no intervention may remove a person's ability to decline** (CANON §5).

Every internal chapter of Volume Two is written for an operator working on themselves, and every external chapter is written for an engineer working on a circumstance that other people are free to enter and leave. Those are different licences, and the correspondence does not transfer the licence with the equation. A technique that is legitimate self-discipline internally becomes coercion externally by the simple substitution of a second person, and the substitution is exactly what the correspondence otherwise licenses.

Chapter 49 supplied the measurable form of this boundary: **an obligation network that removes exit is not a trust system but a capture system, and the two are distinguished by rho, which the first has and the second does not.** The ethical rule and the engineering rule agree here,

which is fortunate but not necessary, and where they disagree the ethical rule wins without argument.

What survives all nine Something does, and it is worth stating precisely rather than defensively.

The recursion transfers. `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` holds at both scopes, and every theorem about it holds at both scopes, because none of them mentions a substrate. §50.13 demonstrates the shared trajectory to within two per cent.

The geometry transfers. `N_eff = 1 + (N−1)u` and `S = ln N_eff` are arithmetic on a redundancy audit, and redundancy is scope-independent. §50.4.

The ceiling transfers, and produces a new result in doing so. `F <= ln N` at both scopes forces `M <= ln N` at internal scope, which is Ashby's requisite variety in first-person form (§50.5). That result was not available at either scope alone.

The diagnostic questions transfer. *Is the maintenance term eating the ceiling? Is the frame outrunning the intake? Is the interaction leaving anything behind?* These are the same three questions with different nouns, and asking them of a person and of a department in the same session, with the same arithmetic, is what this chapter is for.

What does not transfer: anything involving agency of the parts. Exit, heterogeneity of interests, aggregation of preferences, reciprocity, sanction, rivalry of the store. **The correspondence is a correspondence of dynamics and geometry, not of agency.** State it that way and it is defensible. State it as "a company is a mind" and it is a slogan that will be wrong in precisely the six places listed above, all of which are the places where somebody can walk out.

50.15 What would falsify the correspondence

FVP-6 and FVP-7, applied to the thesis of this chapter rather than to a formula. **A framework that cannot fail cannot be trusted, and this is how this one fails.**

Test 1 — cross-scope parameter recovery. Take one operationalisation of `delta` — the two-time zero-input estimator — and apply it to twenty individuals learning a skill and twenty teams maintaining a practice. Preregister. **The correspondence predicts that after normalisation by period length the two `delta` distributions occupy the same functional role**, meaning that `Q_inf = (eta S − C)/delta` predicts held-out `Q` equally well at both scopes with no scope-specific correction term. *If a scope-specific correction term is needed and does not shrink with sample size, the correspondence is false as an equality of role and must be demoted to an analogy of form.*

Test 2 — the shape test. Fit the trajectory shape, not the levels. **The correspondence predicts that normalised trajectories at the two scopes are the same curve** (§50.13). *If internal trajectories are systematically non-exponential in normalised time while external ones are exponential, or vice versa, the shared-dynamics claim fails and the whole chapter reduces to a notational convenience.*

Test 3 — the four baselines, both scopes. Paper E §5.3's template applies unchanged: persistence, direct-input, matched autoregressive, and **shuffled-input control**. *If shuffled U predicts Q as well as real U at either scope, the causal story at that scope is a trend.*

Test 4 — the breakdown predictions. §50.14 is not only a list of caveats; it is a list of predictions. **B2 predicts that a heterogeneous group's fitted `delta` drifts downward over an unmaintained interval while a homogeneous one's does not.** That is directly testable and it is the sharpest available discriminator between the mean-field model and the real one. *If fitted `delta` does not drift in heterogeneous groups, the mean-field caution is wrong, and if it drifts in individuals too, then B2 is not a disanalogy at all and the correspondence is stronger than this chapter claims.*

Test 5 — reproducibility across labs. FVP-7. Same operationalisation, different investigators, different populations. **This gate is the one that decides the hard-science question and it is the one most often skipped.** Nothing in Volume Two has passed it.

The Specification

S50.1 — Scope declared. Every application of the storage law shall declare its scope (internal or external), its period length, and its unit for `Q`, before any parameter is estimated.

S50.2 — Period stated with every rate. `delta`, `eta`, `C` and `h` shall never be reported without their period. Cross-scope comparison shall use `delta_T = 1 − (1−delta)^T` (B7).

S50.3 — Distribution, not scalar, at external scope. `delta`, `eta`, `C` and `u` shall be reported with their across-member distribution at external scope. A scalar alone is a mean-field claim and shall be labelled as one (B2).

S50.4 — Member headroom computed. Any external application reporting `h > 0` shall also report the distribution of `h_i` and name the members below zero (B1). **A body may not be described as sustainable while a quarter of its members are not.**

S50.5 — Rivalry declared. Before any prediction about group size, the analyst shall state whether `Q` is rival or non-rival and justify it (B8).

S50.6 — Cyclicity tested before `D/P`. Group-scope `D/P` shall not be computed until the group's collective ordering has been tested for intransitivity over the relevant option set (B3).

S50.7 — No inward transfer of agency machinery. `rho`, obligation ledgers, sanctions and verification budgets shall not be given internal operationalisations (B4). Their internal use is metaphorical and shall be labelled.

S50.8 — Per-scope validation. An instrument validated at one scope shall be re-validated at the other before use. FVP-7 is per scope (B5).

S50.9 — Licence not transferred. Techniques licensed for self-application shall not be applied to another person on the strength of the correspondence. CANON §5.4 governs (B9).

S50.10 — The thesis carries its badge. In any document quoting this correspondence, it shall be identified as the author's thesis, with §50.14 cited, and not as a proven result. The proven result is the scale-freeness of the *form*, and it is a much smaller claim.

The Failure Modes

FM-50.1 — The Slogan. *Signature:* "a company is just a big brain," or "your mind is a little society," offered as an explanation rather than as a model with limits. *Damages:* everything downstream, because it licenses transfers through all nine breakdowns at once. *Measurement:* ask which of the nine breakdowns the speaker has checked. The usual answer is none. *Fix:* S50.10. State the correspondence as dynamics-and-geometry only.

FM-50.2 — Unnormalised Period Comparison. *Signature:* a table comparing an individual's `delta` to a department's `delta` with no period column. *Damages:* conclusions wrong by up to two orders of magnitude (B7: 90.1x on a realistic pair). *Measurement:* look for the period column. If absent, the comparison is void. *Fix:* S50.2.

FM-50.3 — Mean-Field Complacency. *Signature:* one `delta` for a group with obviously heterogeneous members; a decay forecast that keeps being wrong in the same direction. *Damages:* forecast error rising to 92 per cent over twelve periods (B2). *Measurement:* fit `delta` on early data and on late data separately. Divergence is the diagnosis. *Fix:* S50.3, and compartment the model.

FM-50.4 — Aggregate u Complacency. *Signature:* a diverse-on-paper body reporting a healthy `u` with no distribution. *Damages:* invisible fragility. Configuration B in B2 loses 77 per cent of its surplus to the departure of three of nine against configuration A's 21 per cent, and reports identically beforehand. *Measurement:* compute u with each member removed in turn. The spread of that vector is the fragility. *Fix:* S50.3. This computation costs nine subtractions and should be standing practice.

FM-50.5 — The Group Mind That Cannot Decide. *Signature:* a body that reverses itself in a cycle, with each reversal individually well-argued, and members who each believe the others are being inconsistent. *Damages:* `D` is undefined; every doctrine-to-data diagnostic returns nonsense. *Measurement:* elicit member orderings and check pairwise majorities for a cycle (B3). *Fix:* S50.6. Restrict the option set, or change the aggregation rule, and note that Arrow guarantees you will give up one of the four conditions whichever you choose.

FM-50.6 — The Sustainable Body With Departing Members. *Signature:* positive aggregate headroom, rising attrition, and genuine bewilderment at the attrition. *Damages:* loses the highest-`c_i` members first, who are frequently the highest-u members (B1). *Measurement:* the `h_i` table. It takes an afternoon. *Fix:* S50.4. Reduce `c_i` for the members below zero, or accept the loss with its `ln N` price computed (Chapter 49, FM-49.9).

FM-50.7 — Inward Transfer of the Sanction Ladder. *Signature:* a self-improvement scheme built on internal contracts, penalties and rewards, with a homunculus doing the enforcing. *Damages:* imports a model whose central parameter, `rho`, has no internal referent (B4). *Measurement:* ask what the flow matrix is between which agents. There is no answer. *Fix:* S50.7. The internal input side is the Moment equation, not a negotiation.

FM-50.8 — Introspective Validation. *Signature:* an internal instrument validated by asking the subject whether it seems right. *Damages:* the error in introspection is correlated with the quantity being measured, so agreement is evidence of nothing (B5). *Measurement:* check whether the validation criterion is independent of the instrument. Usually it is not. *Fix:* S50.8, and use behavioural

criteria per Chapter 37.

Worked Example — One Clinician and Her Unit, Measured With One Instrument

A hospital department. The engineer is asked to look at a senior clinician who reports being "completely full" and, separately, at the six-person unit she leads, which is described by its director as "busy but not going anywhere." **The whole point of the exercise is that these are two readings from the same instrument, and the point of §50.14 is the four places where they must be read differently.**

1. Scope declarations (S50.1, S50.2)

	INTERNAL	EXTERNAL
subject	one clinician	the six-person unit
period	1 week	1 quarter
Q unit	retained-capability units (Ch 40)	retained-practice units (Ch 23)
rivalry (S50.5)	not applicable	**non-rival** — the store is shared protocol

2. Baseline, internal

N = 9 distinguishable clinical framings, from a free-sort with test-retest agreement 0.71. u = 0.31 from an independent-prediction round on eight recent cases.

```
S     = ln(1 + 8 * 0.31) = 1.2470          ln N = 2.1972,  so 56.8% of ceiling
eta   = 0.22    (dQ_attributable / M_input over eleven weeks)
C     = 0.42    (time-on-maintenance sampling)
delta = 0.035   (two-time zero-input estimator over a four-week leave)

eta*S = 0.2743        Q_inf = (0.2743 - 0.42)/0.035 = -4.1615
h at Q = 6.0: 0.22 * 2.1972 - 0.035 * 6.0 - 0.42 = 0.4834 - 0.2100 - 0.4200 = -0.1466
```

Negative steady state and negative headroom. Per Theorem 3.5 this is the theorem stating its own failure mode: unsustainable in this model. "Completely full" is h < 0, measured.

3. Baseline, external

N = 6 distinct domains held by the unit, from the capability inventory with the independent-action verification. u = 0.29 from the pre-discussion prediction round.

```
S     = ln(1 + 5 * 0.29) = 0.8961          ln N = 1.7918,  so 50.0% of ceiling
eta   = 0.24
C     = 0.44
```

```
delta_int = 0.05,    turnover tau = 0.10    ->    delta_total = 1 -
(0.95)(0.90) = 0.1450

eta*S = 0.2151        Q_inf = (0.2151 - 0.44)/0.1450 = -1.5513
h at Q = 2.8: 0.24 * 1.7918 - 0.1450 * 2.8 - 0.44 = 0.4300 - 0.4060 - 0.4400 = -
0.4160
```

The same three faults, at both scopes, in the same order: u at roughly half the ceiling, `eta` low, `C` high enough to exceed the ceiling by itself. The clinician and her unit are not two problems; they are one diagnostic profile read twice.

4. The four places the readings must diverge (§50.14)

B7, period. `delta = 0.035` weekly and `delta = 0.145` quarterly are *not* comparable as printed. The clinician's weekly rate over a 13-week quarter is `1 - (0.965)^13 = 0.3707`. **Her personal store decays two and a half times faster per quarter than her unit's does**, which is the opposite of the impression the raw numbers give. This reading would have been inverted without S50.2.

B1 and B4, exit and turnover. The unit's `delta` contains `tau = 0.10` — a term the clinician's does not have and cannot have. Of the unit's total loss, `0.10/0.1450 = 69 per cent` is people leaving. **No amount of documentation addresses that term**, and no internal analogue exists. The interventions therefore cannot be the same interventions.

B2, heterogeneity. Per S50.3, the unit's u was recomputed with each member removed. Removing member 4 dropped u from 0.29 to 0.19; removing any other dropped it by at most 0.02. **One of six people is carrying a third of the unit's remaining non-parallelism, and member 4's `h_i` was the second lowest in the unit.** The scalar `u = 0.29` concealed both facts. No corresponding computation exists internally, because the clinician's framings cannot resign.

B9, licence. The clinician's own `delta`-raising interventions — deliberately letting obsolete protocols lapse — she may apply to herself without further permission. The unit's equivalent requires the members' agreement, and the intervention list below is written so that every member can decline any part of it.

5. Intervention and post-measurement

Internal (weeks 1–20): a second genuinely different framing developed and used on every case (raising u to 0.62); a weekly written case note with a decision line (raising `eta` to 0.34); three standing commitments sunset (lowering `C` to 0.30).

```
S     = ln(1 + 8 * 0.62) = 1.7851       81.2% of ceiling
eta*S = 0.6069                        Q_inf = (0.6069 - 0.30)/0.035 = +8.7693
h at Q = 6.0 = 0.34 * 2.1972 - 0.035 * 6.0 - 0.30 = +0.2371
tau = 28.6 weeks;  half-life = 19.46 weeks
```

External (quarters 1–6): a rotating pre-discussion prediction round (raising u to 0.60); a decision log with named owners (raising `eta` to 0.37); two recurring meetings deleted (lowering `C` to 0.31); and — the term with no internal counterpart — a deliberate retention intervention on the two members whose `h_i` was negative, taking `tau` from 0.10 to 0.04.

```
delta_total = 1 − (0.95)(0.96) = 0.0880
S     = ln(1 + 5 ∗ 0.60) = 1.3863       77.4% of ceiling
eta∗S = 0.5129                      Q_inf = (0.5129 − 0.31)/0.0880 = +2.3060
h at Q = 2.8 = 0.37 ∗ 1.7918 − 0.0880 ∗ 2.8 − 0.31 = +0.1066
tau = 11.36 quarters;  half-life = 7.52 quarters
```

6. The shape test (§50.13, Test 2)

Both trajectories, normalised by their own time constants, from their own starting points:

t / tau	clinician, (Q−Q_0)/(Q_inf−Q_0)	unit, (Q−Q_0)/(Q_inf−Q_0)	1 − e^−k
0.5	0.3927	0.4246	0.3935
1.0	0.6441	0.6370	0.6321
2.0	0.8688	0.8798	0.8647
3.0	0.9533	0.9564	0.9502

Same curve. The clinician's time constant is 28.6 weeks — 2.20 quarters. The unit's is 11.36 quarters. **The unit is five times slower and arrives in the same shape**, which is Banach at both scopes and is the single most useful practical output of this chapter: *you can tell a person how long their change will take by measuring their `delta`, and you can tell a department the same way, and the arithmetic is identical.*

7. `D/P` at both scopes

subject	D	P	r	|ln r|	band
clinician, before	2.40	1.00	2.400	0.8755	fault
clinician, after	1.15	1.00	1.150	0.1398	drift
unit, before	2.90	1.00	2.900	1.0647	**severe**
unit, after	1.25	1.00	1.250	0.2231	drift

Both were over-framed; the unit worse than the clinician. Per S50.6 the unit's option orderings were tested for cyclicity first and were transitive, so `D` was well defined and the diagnostic applied. **Had the test failed, the unit's `r` would have been reported as undefined rather than as 2.90**, and the whole doctrine-to-data intervention would have been the wrong instrument.

8. Falsification, declared in advance

1. If normalised trajectories at the two scopes diverge by more than five per cent at `t/tau` = 1, Test 2 has failed here and the shared-dynamics claim is not supported in this system.
2. If the unit's fitted `delta` drifts downward over the unmaintained interval while the clinician's does not, B2 is confirmed as a genuine scope disanalogy. If neither drifts, the mean-field caution is overstated in this system and should be reported as such.
3. Shuffled-input control at both scopes, mandatory. If permuted `U_t` predicts `Q` as well as the real sequence at either scope, that scope's causal story is a trend.

4. If member 4 leaves and the unit's measured u does not fall by approximately the predicted 0.10, the fragility computation of §50.14 B2 is wrong in this system.

9. The rule, extracted

One instrument, two scopes, one diagnostic profile, and four places where the readings had to be handled differently: the period normalisation, the turnover term, the member-level distribution, and the licence.

The correspondence earned its keep by letting one engineer diagnose a person and a department in the same afternoon with the same arithmetic. It would have destroyed the diagnosis by letting the same engineer treat the department's turnover as forgetting, its dissenting member as a framing, and its members' consent as already given.

A mind is not a small society and a society is not a large mind. They are two systems running the same recursion over different substrates, and the difference between the substrates is exactly this: the parts of a society can leave. Everything the correspondence gets right follows from the shared recursion. Everything it gets wrong follows from that one sentence.

Chapter 51 . The Interpersonal Formulary

Chapter 45 collected every personal formula in one place. This chapter does the same for the interpersonal scope: every formula developed in Part IX, plus the group-scope forms of the Volume One core, in one lookup reference with a fixed entry format.

This chapter is written to be opened, not read through. You arrive with a symptom or a parameter and you leave with a statement, a derivation, a validation verdict, a measurement procedure, a computed example, and the ways the formula fails. Every entry carries the same seven fields in the same order so the chapter can be scanned in a column.

The armature is adapted, as CANON §2 permits for a reference chapter. `The Law` states what makes a formulary possible. `The Engineering` **is** the formulary. `The Specification` governs its use and extension. `The Failure Modes` names the engineer's faults rather than the system's. `The Worked Example` runs one diagnosis through eleven entries end to end.

The Law

51.0.1 What is proven, quoted exactly

```
F(x,y) = f(u(x,y)),   u = 1 − (x.y)^2 = ||x ^ y||^2  in [0,1]
f(u)  = ln(1 + (N−1)u)                          Paper A, Thm 2.1 (unique under S1–S4)
N_eff  = 1 + (N−1)u                                          effective count
0 <= F(u) <= ln N                                            Paper D, Cor 2.2

u = u_cross + u_div,  u_cross = beta^2,  u_div = alpha^2 sin^2 gamma    Paper A, Thm 4.1
max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)

Q_{t+1} = (1−delta)Q_t + eta S_t − C_t                       Paper E, Section 3
Q_t  =  (1−delta)^t  Q_0  +  sum_{k=0}^{t−1}  (1−delta)^{t−1−k}(eta  S_k  −  C_k)   Thm 3.1
Q_{t+1} > Q_t  <=>  eta S_t > delta Q_t + C_t                Cor 3.2
Q_t <= max{Q_0, eta ln N / delta}                            Thm 3.3
C_t <= eta ln N − delta Q_t   for a nondecreasing step       Cor 3.4
Q_inf = (eta F(u∗) − c∗) / delta                             Thm 3.5
h_t = eta ln N − delta Q_t − C_t                             Section 6.4
r_max = eta ln N − delta Q_t − C_t                           Cor 4.2
```

Derived, not axioms: monotonicity (3.1), positivity (3.2), strict concavity (3.3), subadditivity on u (3.4), Lipschitz constant N−1 (4.2), extremality (4.3). **S3, the affine effective count, is the design axiom: chosen, not derived.** Every entry using `N_eff` inherits it.

51.0.2 The seven-field entry format

Field	What it holds
Statement	the formula, in a code fence, with every symbol defined
Derivation	where it comes from: theorem, definition, or postulate — badged
FVP	all seven gates of CANON Addendum II §6, each with a verdict
Protocol	the numbered measurement procedure
Worked number	one computed instance, real arithmetic
Failure modes	the named ways it goes wrong, with signatures
See	the chapters that develop it

Gate names, used in short form in every entry: **FVP-1** dimensions, **FVP-2** scale type, **FVP-3** limiting cases, **FVP-4** fixed points, **FVP-5** identifiability, **FVP-6** falsification, **FVP-7** reproducibility.

The FVP row is mandatory and no entry appears without it (CANON Addendum II §6). Verdicts are one of: **PASS**, **PASS (conditional)** with the condition stated, **FAIL** with the consequence stated, or **OPEN** where the gate has not been run. **OPEN is not a failure and it is not a pass; it is a debt, and it is recorded as one.** Gate 7 is OPEN for twenty-four of the twenty-five entries below and PARTIAL for the twenty-fifth, because nothing in Volume Two has been through a cross-lab reproducibility trial, and pretending otherwise would be the exact fault Chapter 52 is written against.

51.0.3 Honesty clauses governing all twenty-five entries

- The surplus functional as deployed is a **lower bound** on the ensemble's information. Every price in this chapter is a floor.
- The triangle inequality **fails** for geometric triples (Paper A §5). Subadditivity holds on u only; three-party chains do not compose metrically.
- Paper A is pure mathematics and makes **no empirical or metaphysical claim**. Surplus is not identified with Shannon entropy or mutual information (Paper D). **No thermodynamic negentropy law is proved** (Paper E).
- Where an entry cites an external result, the result is real and the **transfer into this parameterisation is a modelling assumption**, badged in the entry.
- **CANON §5.4 governs absolutely: no formula here licenses an intervention that removes a person's ability to decline.** IF-21, IF-22 and IF-24 carry the constraint explicitly.

Two standing verdicts, in force across multiple entries Two findings arrived after these entries were drafted and cut across several of them. They are stated here once and govern wherever they apply.

STANDING VERDICT 1 — every `M`-derived entry is DEMOTED on FVP-2. Chapters 52 and 53 run the protocol on the Moment Equation and return **DEMOTE on gate 2**: `D` and `P` are not instrumented on a common ratio scale with a shared true zero, and division without a true zero produces a number with no referent. This governs **IF-14, IF-15, IF-16 and IF-17** — every entry whose value passes through `M` or through `r = D/P`. Chapter 47 §47.2's channel-counting argument

is corrected there: a count supplies a zero and a unit inside one declared encoder and supplies no linkage between encoders, so the magnitude of `r` moves when the encoder moves.

> **WITHDRAWN as measured quantities:** the magnitude of `D/P`; `Delta = |ln r|` as a number; and the quantitative form of the join from the moment into the storage law. The worked numbers under those four entries are **illustrative constructions demonstrating the arithmetic of each measure.**
>
> **IN FORCE, untouched:** the limiting cases; the `D = P` fixed point; the ordering *over-framed* / *calibrated* / *under-framed*; and the **sign** of `ln r`, which is the operationally decisive part because the two faults take opposite corrections. **What is demoted is magnitude, not structure.**

The route out is constructive, specific and priced. Chapter 54, **Theorem 54.3**, given badged Postulate SO1: measure `D` and `P` as Rasch measures on one common linked scale and `Delta = |ln(D/P)| = |theta_D − theta_P|` becomes genuinely ratio-scaled, Rasch separation supplying the true zero that division requires. §54.12 prices it at **approximately 123 items per term to establish a half-logit distortion at 95 per cent**, and about 31 items for a full logit. At group scope the items are decisions with later-observable answers, so **a body that already keeps a decision log is already accumulating the item bank** — the one respect in which the interpersonal version of this study is cheaper than the personal one. **A formulary that demotes four of its own entries, names exactly what survives, and costs the study that would restore the rest is the protocol working. It is not an apology.**

STANDING VERDICT 2 — the shuffled-input control has three measured failure modes, and every FVP-6 row in this chapter inherits all three. The four-baseline suite of Paper E §5.3 is invoked throughout. Chapters 55 and 81 measured what it actually catches, by simulation rather than assertion:

#	Failure mode	Measured
1	**Anti-conservative as ordinarily run.** Against nominal `alpha = 0.05` over 500 null trials: iid permutation false-rejected at **0.108, 2.2x nominal**; circular shift 0.120; block permutation 0.060, near nominal	Ch 55
2	**Blind to repackaging.** A filtered repackaging passed all four baselines including the shuffle at `p = 0.0005`, and was caught **only by the structural ceiling test** (Paper E Cor 4.2) — its single-step jump was **38.5x** the theoretical maximum	Ch 25, Ch 55

#	Failure mode	Measured
3	**Blind to shared slow drivers.** Across 4,000 trials the shuffle detected the clean case 48.17% of the time and **passed 98.35% of the time against a confound driving both series**	Ch 79–81

THE RULE, and it supersedes any earlier statement of it: the shuffle is a weak, mis-calibrated empirical audit. Repackaging is caught only by the structural ceiling test; confounding by a shared slow driver is caught only by design controls — blinding, channel separation, prospective sampling, a yoked control. No surrogate catches either. Calibrate the surrogate by simulation on your own design before using it, and **design controls, not more surrogates.** Chapter 55 specifies both.

51.0.4 Two indexes

By parameter. Which entries move which quantity:

Parameter	Entries
`N`	IF-02, IF-04, IF-05, IF-22
`u`	IF-01, IF-02, IF-03, IF-05
`S` / `F`	IF-01, IF-02, IF-03, IF-04
`Q`	IF-06, IF-07, IF-08
`delta`	IF-11, IF-18, IF-19, IF-25
`eta`	IF-20, IF-21
`C`	IF-12, IF-13, IF-19
`h`	IF-09, IF-10, IF-22
`M`, `D/P`	IF-14, IF-15, IF-16, IF-17
thresholds	IF-23, IF-24

By symptom. What you arrived with:

Symptom	Start at
“lots of people, no new ideas”	IF-02, then IF-05
“we agree things and start over every quarter”	IF-08, then IF-20
“meetings are necessary and useless”	IF-12, then IF-13
“we’re busy but nothing accumulates”	IF-20, IF-21
“we keep being surprised by the same thing”	IF-16, IF-17
“good people keep leaving”	IF-22, IF-11
“the reorganisation didn’t help”	IF-09, IF-10
“one person left and everything got worse”	IF-05
“we cannot get the change to stick”	IF-24

Symptom	Start at
“the numbers from the team and the person disagree”	IF-25

The Engineering — the formulary

Class A — Geometry of difference

IF-01 . The pair surplus Statement.

```
S_2 = ln(1 + u),     u = 1 − (x . y)^2  in [0,1]
ceiling: ln 2 = 0.693147
```

x, y unit vectors encoding the two parties’ framings; u their non-parallelism; S_2 the surplus of the dyad in nats per interaction.

Derivation. Paper A Theorem 2.1 at `N = 2`: `f(u) = ln(1 + (2−1)u) = ln(1+u)`. Not a special case requiring separate proof — the theorem holds for all `N >= 1`. **Extremality (Thm 4.3): `S_2 = 0` if and only if the framings are parallel.** This is the formal content of the canon statement that *the wedding of like with like is a non-event*.

FVP.

Gate	Verdict
FVP-1	**PASS.** u dimensionless by construction (squared cosine); S_2 in nats
FVP-2	**PASS.** u is ratio-scaled with a true zero (parallel) and a true ceiling
FVP-3	**PASS.** `u -> 0` gives `S_2 -> 0` (identical partners generate nothing); `u -> 1` gives `ln 2`, the hard ceiling. Both are real, nameable states
FVP-4	**PASS.** `u = 0` is the degenerate fixed point; no interaction changes it
FVP-5	**PASS (conditional)** on the encoding. u is recoverable from independent-prediction disagreement; it is **not** recoverable from self-reported difference
FVP-6	**PASS (conditional).** Preregister the encoding, freeze, predict held-out joint output. Killed if pairs at `u = 0.1` and `u = 0.8` produce indistinguishable output
FVP-7	**OPEN.** No cross-lab recovery of dyadic u has been run

Protocol. 1. Choose the framing space and justify the basis (Paper B validity conditions govern). 2. Each party independently predicts k >= 8 observables in the domain, without conferring. 3. Encode predictions as unit vectors; compute u = 1 − (x.y)^2. 4. Repeat after one period. Report test-retest; below 0.6 the estimate is not a measurement.

Worked number.

u	0.00	0.10	0.25	0.50	0.75	0.90	1.00
S_2	0.0000	0.0953	0.2231	0.4055	0.5596	0.6419	0.6931
% of ln 2	0.0	13.8	32.2	58.5	80.7	92.6	100.0

Half the available surplus of a pair is bought by the first quarter of the way to orthogonality (u = 0.25 gives 32.2%; u = 0.50 gives 58.5%). Strict concavity (Thm 3.3), arriving as a design fact.

Failure modes. *Similarity mistaken for compatibility* — a pair chosen for agreement has S_2 = 0 exactly. *Encoded difference* — vocabulary differences encode as u > 0 while predicting identically; the independent-prediction step excludes them. *Conscripted difference* — an instructed disagreement produces u with no informational content, and is prohibited by CANON §5.4 where the partner cannot decline the role.

See: Chapter 46; Paper A Thm 2.1, 4.3.

IF-02 . The effective count and the group surplus Statement.

```
N_eff = 1 + (N − 1)u
S     = ln N_eff = ln(1 + (N−1)u)
0 <= S <= ln N
```

N the count of distinct domains held by the group (not headcount); u the mean pairwise non-parallelism; S the group's surplus per interaction, in nats.

Derivation. Paper A Theorem 2.1, with the effective count coming from axiom **S3**, the affine requirement on g(u) = e^{f(u)}. **S3 is the design axiom — it is chosen, not derived — and this entry must not be used without saying so.** What S3 buys is that surplus is the log of a headcount *after the headcount has been audited for redundancy*.

FVP.

Gate	Verdict
FVP-1	**PASS.** N a count, u dimensionless, S in nats
FVP-2	**PASS (conditional).** N is ratio-scaled if domains are enumerable by the independent-action test; if N is headcount, the entry is being misused
FVP-3	**PASS.** u −> 0 gives N_eff −> 1, S −> 0: the group is one instrument. u −> 1 gives S −> ln N

Gate	Verdict
FVP-4	**PASS.** `u = 0` is the collapse fixed point
FVP-5	**PASS (conditional).** `N` and `u` are separately identifiable only if the domain inventory is verified independently of the prediction round; otherwise they trade off
FVP-6	**PASS (conditional).** Shuffled-input control mandatory (Paper E §5.3)
FVP-7	**OPEN**

Protocol. 1. **Capability inventory.** Enumerate domains in which at least one member acts without external consultation, each verified by an observed independent instance. That count is `N`. 2. **Pre-discussion prediction round.** All members predict independently before any discussion. 3. `u` = mean pairwise `1 − (x_i . x_j)^2` over all `C(N,2)` pairs. 4. Compute `N_eff` and `S`. **Never measure `N_eff` directly** — it is derived. 5. Report `S / ln N`, the realised fraction of the ceiling, alongside `S`.

Worked number. `N = 12`, `ln 12 = 2.4849`:

u	N_eff	S	% of ceiling
0.00	1.000	0.0000	0.0
0.10	2.100	0.7419	29.9
0.20	3.200	1.1632	46.8
0.35	4.850	1.5790	63.5
0.50	6.500	1.8718	75.3
0.65	8.150	2.0980	84.4
0.80	9.800	2.2824	91.8
1.00	12.000	2.4849	100.0

Twelve people at `u = 0.10` are worth 2.1 people. Not “somewhat redundant” — 2.1.

Failure modes. *Headcount substituted for `N`* — the single most common error in this book’s subject and the reason Chapter 11 exists. *Reporting S without `S / ln N`* — half a diagnosis; the ceiling and the realised fraction are different levers. *Reporting `N_eff` from a survey* — it is derived and cannot be asked about.

See: Chapters 7, 8, 11; Paper A Thm 2.1.

IF-03 . The two-source decomposition Statement.

```
u = u_cross + u_div,    u_cross = beta^2,   u_div = alpha^2 sin^2(gamma)
max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

`u_cross` between-category non-parallelism; `u_div` within-category non-parallelism.

Derivation. Paper A Theorem 4.1, for a pure block vector. **The upper bound is subadditivity on the parameter u (Thm 3.4) and is not a triangle inequality** — Paper A §5 shows the triangle inequality fails for geometric triples, and this entry must not be chained.

FVP.

Gate	Verdict
FVP-1	**PASS.** All three terms dimensionless
FVP-2	**PASS.** Additive on u, which is ratio-scaled
FVP-3	**PASS.** `u_div -> 0` gives pure cross-category difference; `u_cross -> 0` gives pure within-category spread. Both are real organisational states
FVP-4	**PASS.** `u_cross = u_div` is the balanced-block fixed point
FVP-5	**FAIL as stated.** Many `(u_cross, u_div)` pairs give the same u, and the aggregate does not identify the split. **The split must be measured from the block assignment, not inferred from u**
FVP-6	**PASS (conditional)** on an independently specified block structure
FVP-7	**OPEN**

Protocol. 1. Declare the category partition **before** measuring, from an external criterion (function, site, discipline). A partition chosen after seeing the data is not a partition, it is a fit. 2. Compute u within each block and across blocks separately. 3. Verify the bound. **A violation means the block assignment is wrong, not that the theorem is.**

Worked number. `N = 12`, `u_cross = 0.24`, `u_div = 0.16`, so `u = 0.40`:

```
f(u)       = ln(1 + 11 * 0.40) = 1.6864
f(u_cross) = ln(1 + 11 * 0.24) = 1.2920
f(u_div)   = ln(1 + 11 * 0.16) = 1.0152
bound:  1.2920 <= 1.6864 <= 2.3072      satisfied
subadditivity slack = 2.3072 - 1.6864 = 0.6208
```

The slack of 0.6208 nats is the amount by which naive addition of the two sources would overstate the surplus — **37 per cent of the true value.** Any report that sums category diversity and within-category diversity is making that error.

Failure modes. *Chaining across three parties* — prohibited; the triangle inequality fails. *Post-hoc partition* — the block structure must be declared first. *Summing the sources* — see the slack above.

See: Chapters 10, 11; Paper A Thm 4.1, §5.

IF-04 . Requisite variety at group scope Statement.

```
required:  N_eff >= N_env
hence:     u >= (N_env - 1) / (N - 1)
feasible only if N >= N_env, since u <= 1
```

N_env the number of distinguishable states the environment presents that the group must respond to differently.

Derivation. Ashby's Law of Requisite Variety (1956): only variety absorbs variety. **The transfer into this parameterisation is a modelling assumption, badged.** What is proven here is only the algebra: setting `1 + (N−1)u >= N_env` and solving for u.

FVP.

Gate	Verdict
FVP-1	**PASS.** Both sides counts
FVP-2	**PASS (conditional).** Requires N_env to be a genuine enumeration, which is the hard part
FVP-3	**PASS.** `N_env = 1` requires nothing; `N_env > N` is infeasible at any u
FVP-4	**PASS.** `N_eff = N_env` is the exact-sufficiency fixed point
FVP-5	**PASS (conditional)** on independent enumeration of N_env
FVP-6	**PASS.** Killed if groups with `N_eff < N_env` respond correctly to the full state space at rates indistinguishable from those with `N_eff >= N_env`
FVP-7	**OPEN**

Protocol. Enumerate the environment states requiring **distinct** responses, merging any two whose correct response is identical. Compute `N_eff` per IF-02. If `N_eff < N_env`, the shortfall is the number of states the group must handle by guessing.

Worked number.

N_env	required u at N = 8	at N = 12	at N = 20
4	0.4286	0.2727	0.1579
8	1.0000	0.6364	0.3684
16	infeasible	infeasible	0.7895
32	infeasible	infeasible	infeasible

An eight-member body facing an eight-state environment must be perfectly orthogonal — `u = 1.0` — to be adequate, which is not achievable in practice. That is the arithmetic reason understaffed bodies fail on variety rather than on effort.

Failure modes. *Inflated N_env* — merging states with identical correct responses is the discipline; without it every environment looks infinite. *Treating the requirement as sufficient* — `N_eff >= N_env` is necessary, not sufficient; the group must also convert (`eta`) and afford (`C`).

See: Chapters 8, 18; Ashby 1956.

IF-05 . Departure fragility: targeted versus random removal Statement.

```
model:  u_ij = 1 − c_i c_j,  c_i = cos^2(th_i)   [shared direction + private direction]
mean u = 1 − [ (sum c)^2 − sum c^2 ] / (2 ∗ C(N,2))

fragility = S(N,u) − S(N−1, u after removing the lowest−c member)
```

Derivation. Definition plus Paper A Thm 2.1. **The one-parameter geometry is a modelling choice, badged**: it assumes each member holds a shared direction `e_0` plus one private orthogonal direction, which is the minimal structure that produces heterogeneous `c_i`.

FVP.

Gate	Verdict
FVP-1	**PASS**
FVP-2	**PASS.** `c_i` is a squared cosine, ratio-scaled
FVP-3	**PASS.** All `c_i = 1` gives `u = 0`; all `c_i = 0` gives `u = 1`
FVP-4	**PASS.** Homogeneous `c` is the departure-invariant fixed point
FVP-5	**FAIL for the aggregate.** Distinct `{c_i}` vectors give identical mean u. **The aggregate does not identify the fragility and this is the entry's whole point**
FVP-6	**PASS.** Predicts the u drop for a named departure. Killed if measured post-departure u does not match
FVP-7	**OPEN**

Protocol. 1. Estimate each `c_i` from the member's prediction vector against the group centroid. 2. Compute mean u per the closed form. 3. **Recompute u with each member deleted in turn.** The vector of deltas is the fragility profile. 4. Report the maximum, not the mean.

Worked number. `N = 9`, six members at `c = 0.95` and three at `c = 0.0501`, giving `u = 0.6000` and `S = 1.7578`:

removal	resulting S at `N = 8`	loss
random member (expectation)	1.6434	0.1144
a conformist (`c = 0.95`)	1.7162	0.0416
the most divergent (`c = 0.05`)	**1.4979**	**0.2599**

Targeted removal costs 2.27 times random removal. This is Albert, Jeong & Barabási's result (*Nature*, 2000) — scale-free structures are robust to random failure and fragile to targeted attack — appearing in the surplus geometry rather than in a connectivity graph. **The transfer is a modelling assumption; the arithmetic above is exact.**

And the deeper case: two nine-member groups can have **identical** `u = 0.5999` and `S = 1.7578`, and lose either 1.3862 or **0.3971** of surplus when three members leave, depending on which three. Same aggregate, 3.49 times the damage.

Failure modes. *Scalar u reporting* — conceals the entire profile. *Assuming the loudest member is the most divergent* — divergence is measured by prediction disagreement, not by volume. *Retention policy by seniority* — seniority is uncorrelated with `c_i` and the fragility profile is the correct basis, subject to the constraint that **retention is an offer and never a restriction** (CANON §5.4).

See: Chapters 11, 20, 22; Chapter 50 §50.14 B2.

Class B — Storage and dynamics

IF-06 . The institutional storage law Statement.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = ln(1 + (N−1)U_t)
admissibility:  C_t <= (1 − delta) Q_t + eta S_t   guarantees Q_{t+1} >= 0
one−step growth: Q_{t+1} > Q_t  <=>  eta S_t > delta Q_t + C_t
```

Derivation. Paper E §3. **The mathematics is proven; the mapping of an institution onto a single scalar `Q` is a postulate, badged** (Chapter 27 §27.1 states this in full and it governs here).

FVP.

Gate	Verdict
FVP-1	**PASS.** `Q` in retained-practice units; `delta` dimensionless per period; `eta` in `Q`-units per nat; `C` in `Q`-units per period. **Both sides in `Q`-units**
FVP-2	**PASS (conditional).** Requires `Q` on a ratio scale with a true zero. The substitution test supplies one; a Likert culture survey does not
FVP-3	**PASS.** `delta −> 0` gives unbounded linear growth `Q_0 + t eta ln N` (Thm 3.3) — no fixed point, no arrival. `delta −> 1` gives `Q_{t+1} = eta S_t − C_t`, pure memoryless flow
FVP-4	**PASS.** `Q_inf = (eta S − C)/delta`, entry IF-08

Gate	Verdict
FVP-5	**PASS (conditional).** `eta` and `S` are separately identifiable only if `S` is measured independently of `Q` — otherwise only the product is recoverable
FVP-6	**PASS.** The four baselines of Paper E §5.3: persistence, direct-input, matched autoregressive, **shuffled-input control**
FVP-7	**OPEN**

Protocol. 1. Define the `Q` unit and its zero, by the substitution test: replace the individual and score what survives. 2. Measure `Q` at consecutive period boundaries for at least `3/delta` periods. 3. Measure `S` independently, per IF-02. 4. Fit `delta`, `eta`, `C` on training periods; **freeze**; report held-out prediction against the four baselines.

Worked number. `N = 12`, `u = 0.45`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, `Q_0 = 0`:

`S = ln(1 + 11*0.45) = 1.7834  eta*S = 0.7134  Q_inf = (0.7134-0.55)/0.10 = 1.6336`

t	0	1	2	3	5	8	10	12
Q	0.0000	0.1634	0.3104	0.4427	0.6690	0.9304	1.0640	1.1722

Half-life to the fixed point `ln2 / -ln(1-0.10) = 6.579` periods; 95 per cent settling at 28.43.

Failure modes. *Ordinal `Q`* — a five-point survey cannot enter this equation (FVP-2). *`S` inferred from `Q`* — destroys identifiability. *Declaring success at `t = 3`* — at `t = 3` the system is 27 per cent of the way to its own equilibrium.

See: Chapters 1, 2, 23; Paper E §3.

IF-07 . The exact solution and the contraction Statement.

```
Q_t = (1-delta)^t Q_0 + sum_{k=0}^{t-1} (1-delta)^{t-1-k} (eta S_k - C_k)      Thm 3.1
|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|                                constant input
tau = 1/delta   periods;   t_half = ln2 / -ln(1-delta)
```

Derivation. Paper E Theorem 3.1 by induction. The contraction statement follows because the map is affine with slope `(1-delta)`; **Banach's fixed point theorem** then guarantees a unique fixed point and convergence from any start.

FVP.

Gate	Verdict
FVP-1	**PASS**
FVP-2	**PASS** given IF-06's condition
FVP-3	**PASS.** `delta = 0` breaks the contraction and the fixed point ceases to exist — the transient never decays. This is the tradition regime (Chapter 1)
FVP-4	**PASS.** The fixed point is IF-08
FVP-5	**PASS.** `delta` is recoverable from the decay of the transient alone, under zero input
FVP-6	**PASS.** Predicts the full trajectory shape, not just the endpoint. Killed by a non-geometric approach
FVP-7	**OPEN**

Protocol. Measure `Q` at two times under zero input; `delta = 1 − Q_1/Q_0`. Verify geometry with a third point: `Q_2/Q_1` must equal `Q_1/Q_0` within measurement error, **and if it does not, the system is not single-compartment** (Chapter 50 §50.14 B2).

Worked number.

`delta`	`tau`	half-life	95% at `t`	99% at `t`
0.05	20.00	13.51	58.40	89.78
0.10	10.00	6.58	28.43	43.71
0.20	5.00	3.11	13.43	20.64
0.35	2.86	1.61	6.95	10.69

Failure modes. *Two-point `delta` on a heterogeneous store* — gives a number that is neither compartment's and drifts with when you looked. *Measuring under non-zero input* — the estimator requires zero input and conflates `delta` with `eta S − C` otherwise.

See: Chapters 2, 44; Paper E Thm 3.1.

IF-08 . The group steady state Statement.

`Q_inf = (eta F(u*) − c*) / delta` `delta > 0, constant U == u*, C == c*`
`if eta F(u*) < c*  the equilibrium is NEGATIVE`

Derivation. Paper E Theorem 3.5, by setting `Q_{t+1} = Q_t`. **The theorem states its own failure mode and this entry must reproduce it:** a negative `Q_inf` means the regime is unsustainable, or the linear model must be replaced. It does not mean negative stored order, which is meaningless.

FVP.

Gate	Verdict
FVP-1	**PASS.** Q-units
FVP-2	**PASS** given IF-06's condition
FVP-3	**PASS.** `delta -> 0` gives `Q_inf -> infinity` — correctly, since the fixed point vanishes. `eta F -> c*` gives `Q_inf -> 0`
FVP-4	**PASS.** This *is* the fixed point
FVP-5	**PASS (conditional).** Only the combination `(eta F - c)/delta` is identifiable from steady-state data alone; the three parameters require transient data
FVP-6	**PASS.** Killed if measured long-run Q is not predicted by the fitted parameters out of sample
FVP-7	**OPEN**

Protocol. 1. Confirm the input is approximately constant over at least `3/delta` periods, or the steady state is not the right object. 2. Fit `eta`, `C`, `delta` from the transient (IF-07), then predict `Q_inf`. 3. **Check the sign of `eta F - c` first.** If negative, stop and report unsustainability; do not compute a negative equilibrium and present it as a small number.

Worked number. `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`:

u	S	eta S	Q_inf
0.45	1.7834	0.7134	+1.6336
0.70	2.1633	0.8653	+3.1533

Raising u from 0.45 to 0.70 nearly doubles the equilibrium: **+1.5197 units for a change that costs nothing in headcount.**

Failure modes. *Reporting a negative `Q_inf` as a forecast* — it is a declaration that the model has left its domain. *Applying it to a non-stationary input* — the steady state of a system whose input is trending is not defined by this formula.

See: Chapters 6, 23; Paper E Thm 3.5.

IF-09 . Organisational headroom Statement.

```
h_t = eta ln N - delta Q_t - C_t

h_t < 0  unsustainable in this model. No interaction geometry saves it.  (Cor 3.4)
h_t = 0   maintenance only. Hold, do not build.
h_t > 0  growth available; any requested increment r_t must satisfy r_t <= h_t.
h_t = 0 at  Q_t = (eta ln N - C_t)/delta
```

Derivation. Paper E §6.4. Note that headroom uses `ln N`, the **ceiling**, while the growth condition uses `S = ln N_eff`, the **realised** value. **Half of all misdiagnosis turns on that distinction** (Chapter 27 §27.3): a system can have large positive headroom and shrink every period, because it operates far below its own ceiling.

FVP.

Gate	Verdict
FVP-1	**PASS.** All three terms in Q-units per period
FVP-2	**PASS** given IF-06's condition
FVP-3	**PASS.** `Q -> 0` gives `h -> eta ln N − C`, the greenfield headroom. `Q` large gives `h -> −infinity`: **success consumes its own headroom**
FVP-4	**PASS.** `h = 0` at `Q = (eta ln N − C)/delta`, the ossification point
FVP-5	**PASS (conditional)** on `N` being the domain count and not headcount
FVP-6	**PASS.** Killed if systems with `h < 0` reliably grow
FVP-7	**OPEN**

Protocol. Derive from the four parameters. **Never survey it.** "Do you have capacity" measures perceived h, which is precisely the quantity distorted by the doctrine-to-data fault (IF-16).

Worked number. `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, so `eta ln N = 0.9940`:

Q	0	1	2	3	4	5	7
h	+0.4440	+0.3440	+0.2440	+0.1440	+0.0440	-0.0560	-0.2560

Ossification at `Q = 4.4396`. **The organisation runs out of headroom by succeeding**, and the number at which it happens is computable in advance.

Failure modes. *Using S instead of `ln N`* — understates the ceiling and hides available growth. *Reporting h without `S / ln N`* — half a diagnosis (Chapter 27 §27.3). *Treating `h < 0` as a motivation problem* — Corollary 3.4 states that **no arrangement of people prevents decline** in that regime.

See: Chapters 5, 6, 27.

IF-10 . The target increment and the maximum step Statement.

```
Q_{t+1} − Q_t >= r_t   iff   S_t >= (delta Q_t + C_t + r_t) / eta
feasible only if  delta Q_t + C_t + r_t <= eta ln N
r_max = eta ln N − delta Q_t − C_t = h_t
required u:  u = (e^{S_required} − 1)/(N − 1),  feasible only if u <= 1
```

Derivation. Paper E Proposition 4.1 and Corollary 4.2. **No policy based on interaction geometry alone can guarantee more than `r_max`.**

FVP.

Gate	Verdict
FVP-1	**PASS**
FVP-2	**PASS** given IF-06's condition
FVP-3	**PASS.** `r -> 0` reduces to the one-step growth condition; `r -> h` saturates
FVP-4	**PASS.** `r = h` is the boundary
FVP-5	**PASS (conditional)** on `eta` being estimated independently
FVP-6	**PASS.** Killed by an observed increment exceeding `r_max` with `N` unchanged
FVP-7	**OPEN**

The concavity correction, and it is required before either feasibility test is read. The statement above is written as though `U` were constant every period. It never is. What accumulates is `sum of eta F(U_t)`, so **the operative quantity is the mean of `F(U_t)`, never `F` of the mean `U`.** Because `F` is strictly concave (Paper A, Theorem 3.3), Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )       with equality only if U is constant
```

A body whose geometry varies delivers strictly less surplus than one that sits at its own average every period, so testing feasibility on `F(mean U)` approves increments that cannot be delivered. The gap is the concavity penalty. In the worked office of Chapter 26 §26.7 it is 0.0761 nats — mean `F` of 1.1501 against `F(mean U)` of 1.2262, **a loss of 6.2 per cent of the surplus purely to variability** — and there it was enough to turn a feasible-looking plan into a shortfall. Report `u_req` for legibility; do not test against it. Regularising a variable geometry recovers the penalty for free and is usually the cheapest surplus in the building.

Protocol. 1. State the requested increment `r_t` in `Q`-units. **A target not expressed in `Q`-units cannot be priced and should be refused at this step.** 2. Compute the required `S`, then the required `u`. **Compare the required `S` against the measured mean of `F(U_t)`, not against `F` of the measured mean `U`.** 3. If required `u > 1`, **the target is infeasible at the current `N`** and the only remaining move is to raise `N` (Chapter 8) or lower `C` (Chapter 5).

Worked number. `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, `Q_t = 2.0`, requested `r_t = 0.30`:

```
required S = (0.10*2.0 + 0.55 + 0.30)/0.40 = 2.6250
ln N = 2.4849  ->  required S exceeds the ceiling. INFEASIBLE.
required u = (e^2.6250 - 1)/11 = 1.1641 > 1.  Confirmed infeasible.
r_max = 0.9940 - 0.2000 - 0.5500 = 0.2440
```

The most the body can deliver is 0.2440 and it was asked for 0.3000. The gap is 23 per cent and is not closeable by effort, exhortation, or reorganisation of the same twelve domains.

Failure modes. *Accepting an unpriced target* — the commonest governance failure this entry addresses. *Re-planning instead of re-scoping* — when required `u > 1` the plan is not the problem.

Feasibility tested on F(mean U) — the plan clears the gate on paper and falls short every period; test on mean F(U_t), worth 6.2% in the Ch 26 case.

See: Chapters 6, 26; Paper E Prop 4.1, Cor 4.2.

IF-11 . Turnover composition Statement.

`delta_total = 1 − (1 − delta_int)(1 − tau)`

`delta_int` the intrinsic decay of practice; `tau` the per-period turnover rate. Survivals multiply because the two losses are independent.

Derivation. Definition plus an independence assumption. **The independence is a postulate, badged**, and it is violated in the specific case where departures are caused by the decay (which is the exit feedback of Chapter 50 §50.14 B1) — in that case the composition understates the loss.

FVP.

Gate	Verdict
FVP-1	**PASS.** Both rates dimensionless per period, same period
FVP-2	**PASS.** Both are proportions with true zeros
FVP-3	**PASS.** `tau = 0` recovers `delta_int`; `tau = 1` gives total loss regardless of `delta_int`
FVP-4	**PASS.** `delta_total = delta_int` iff `tau = 0`
FVP-5	**PASS.** `tau` is directly observable from the roster; `delta_int` from a stable-roster interval
FVP-6	**PASS.** Killed if measured `delta_total` in a high-turnover period does not match the composition
FVP-7	**OPEN**

Protocol. 1. Measure `tau` from the roster: departures per period over average headcount. 2. Measure `delta_int` per IF-07 during an interval with no departures. 3. Compose. **Check the independence assumption** by testing whether departures cluster after `Q` declines; if they do, use the composed value as a lower bound only.

Worked number. `delta_int = 0.05`:

`tau`	`delta_total`	`Q_inf` relative to `tau = 0`
0.00	0.0500	100.0%
0.05	0.0975	51.3%
0.10	0.1450	34.5%
0.20	0.2400	20.8%
0.30	0.3350	14.9%

tau	delta_total	Q_inf relative to tau = 0
0.50	0.5250	9.5%

Five per cent turnover halves steady-state stored order. Not reduces — halves. This is the single most under-priced parameter in institutional design.

Failure modes. *Adding the rates* — gives 0.15 where the truth is 0.1450; small here and large at high `tau`. *Attributing the loss to documentation quality* — at `tau = 0.20`, 79 per cent of the total loss is people leaving. *Assuming replacement restores Q* — the replacement restores `N`, not `Q`; `Q` must be rebuilt at rate `eta S`.

See: Chapters 3, 23; Chapter 50 §50.7.

Class C — Cost and structure

IF-12 . The coordination cost law and the crossover Statement.

```
paths(n) = n(n−1)/2                    Brooks 1975
C(n)     = c_0 + c_1 ∗ n(n−1)/2
crossover n∗: the least n with  C(n) > eta ln n
```

Derivation. The path count is exact combinatorics. **The linear-in-paths cost model is a postulate, badged.** The crossover follows because `C` grows quadratically while surplus grows logarithmically.

FVP.

Gate	Verdict
FVP-1	**PASS.** c_1 in Q-units per path per period
FVP-2	**PASS (conditional).** c_1 requires a ratio-scaled cost measure — time, money, or a defined effort unit
FVP-3	**PASS.** `n = 1` gives zero paths; `n –> infinity` gives `C –> infinity` while `eta ln n` grows without bound but far more slowly
FVP-4	**PASS.** n∗ is the crossover, and it is the only fixed point of interest
FVP-5	**PASS.** c_1 is recoverable by regressing observed coordination cost on `n(n−1)/2` across sub-units
FVP-6	**PASS.** Killed if coordination cost is observed to grow linearly in n across a range where the model predicts quadratic growth
FVP-7	**OPEN**

Protocol. Measure coordination cost per period across sub-units of differing size in one unit; regress on n(n−1)/2 to recover c_1, reporting the fit. **A poor fit means the group is not all-to-all connected** — the good case, and it points to IF-13. Then compute n∗.

Worked number. eta = 0.40, c_1 = 0.0009:

n	paths	C	eta ln n	h at delta Q = 0
12	66	0.0594	0.9940	+0.9346
20	190	0.1710	1.1983	+1.0273
30	435	0.3915	1.3605	+0.9690
40	780	0.7020	1.4756	+0.7736
50	1225	1.1025	1.5648	+0.4623
60	1770	1.5930	1.6377	+0.0447
70	2415	2.1735	1.6994	**-0.4741**

Crossover n∗ = 61. Sensitivity to c_1:

c_1	0.0005	0.0010	0.0020	0.0050	0.0100
n∗	85	58	39	23	16

A twenty-fold difference in per-path cost moves the viable group size from 85 to 16. This is why "what is the right team size" has no context-free answer and why c_1 must be measured rather than assumed.

Failure modes. *Assuming a universal team size* — see the sensitivity table. *Measuring c_1 from meeting hours only* — omits the asynchronous coordination cost, which is usually larger. *Adding people to a late project* — Brooks's own result; the added C is immediate and the added S is delayed.

See: Chapters 5, 18; Brooks 1975.

IF-13 . The nesting law Statement.

```
span g, depth d:  N = g^d,  groups = (g^d − 1)/(g − 1),  paths = groups ∗ g(g−
1)/2
paths per member −> g/2  as depth grows
flat coordination cost is O(N^2);  nested is O(N)
```

Derivation. Exact combinatorics on a g-ary tree. **The mapping onto Ostrom's eighth design principle — nested enterprises — is a modelling assumption, badged.** Ostrom (1990) reported nesting as an empirical regularity among long-enduring institutions; the cost arithmetic is this manual's.

FVP.

Gate	Verdict
FVP-1	**PASS.** Counts throughout
FVP-2	**PASS**
FVP-3	**PASS.** `d = 1` recovers the flat case exactly; `g = 2` minimises paths per member at 1.0 and maximises depth
FVP-4	**PASS.** `paths/member = g/2` is the depth-invariant limit
FVP-5	**PASS.** `g` and `d` are directly observable from the org structure
FVP-6	**PASS.** Killed if measured coordination cost in a nested body scales as `N^2`
FVP-7	**OPEN**

Protocol. 1. Map the actual communication structure, not the published chart. Conway's Law (1967) means the two diverge and the actual one is the one that costs. 2. Count paths that carry real traffic. 3. Compare with `groups * g(g−1)/2` for the nominal structure. The excess is the cost of the informal cross-links, and **it is not waste** — see the caveat below.

Worked number.

g	d	N	groups	nested paths	flat paths	ratio
3	3	27	13	39	351	1 : 9
4	3	64	21	126	2,016	1 : 16
5	3	125	31	310	7,750	1 : 25
6	3	216	43	645	23,220	1 : 36
8	3	512	73	2,044	130,816	1 : 64
10	4	10,000	1,111	49,995	49,995,000	1 : 1,000

This is the only entry in the formulary that changes an exponent rather than a coefficient, and the surplus ceiling `eta ln N` is untouched by it, so nesting is a pure reduction of `C`.

The mandatory caveat. A `g^d` tree is exactly the structure Alexander's *A City is Not a Tree* (1965) identifies as wrong: **a strict tree amputates the cross-links that a semilattice preserves, and those links are where u lives.** The figures above are the cost floor of the tree, not a recommendation. A semilattice costs more paths and preserves the overlaps; the correct design is the cheapest semilattice that keeps the cross-links the work needs, and the tree's number is the benchmark it should be compared against.

Failure modes. *Nesting as a tree* — see the caveat. *Nesting the chart without nesting the traffic* — Conway's Law; the cost follows the traffic. *Adding depth without limit* — depth is a delay and a distortion term that appears in IF-16, not in `C`; every tier is an opportunity for doctrine to be reasserted and data to be summarised away.

See: Chapters 5, 21; Ostrom 1990; Alexander 1965; Conway 1967.

IF-14 . The group launch budget Statement.

`L = max(0, M − C)`

`L` the group's capacity to imagine and pursue an alternative to its present arrangement; `M` the group moment (IF-15); `C` the dissipation.

Derivation. Book One Node 00's departure-lounge doctrine, formalised in CANON Addendum II §8 and carried to group scope. **The claim that collective alternatives are generated only from surplus is a postulate, badged.** The `max(0, .)` is the model's statement that there is no negative imagination — below the cost line the quantity is zero, not negative.

FVP.

Gate	Verdict
FVP-1	**PASS (conditional).** `M` and `C` must be in the same unit. Chapter 47 normalises `M` into `[0, ln N]` for exactly this reason
FVP-2	**DEMOTE** (Standing Verdict 1). The subtraction requires an interval scale at minimum; the normalisation into `[0, ln N]` requires ratio. `M` inherits IF-15's demotion, so **L is reportable as a threshold condition — is the moment above the cost line or not — and not as a magnitude.** The sign of `M − C` is the finding; its size is not, until Ch 54 Thm 54.3's instrument exists (~123 items per term for half a logit)
FVP-3	**PASS.** `M <= C` gives `L = 0`: a body with no surplus does not generate alternatives. `C -> 0` gives `L -> M`
FVP-4	**PASS.** `M = C` is the threshold; below it the group is fully occupied maintaining itself
FVP-5	**PASS (conditional)** on independent measurement of `M` and `C`
FVP-6	**PASS.** Predicts that the count of distinct alternatives a body generates per period is monotone in `M − C`. Directly testable and directly killable
FVP-7	**OPEN**

Protocol. Normalise `M` into `[0, ln N]` per Chapter 47; measure `C` in the same unit and period; count distinct alternatives actually generated per period as the dependent variable.

Worked number.

M	C	L	reading
1.20	0.55	0.6500	alternatives available

M	C	L	reading
2.00	0.55	1.4500	ample
0.80	0.90	0.0000	none — fully consumed
0.55	0.55	0.0000	exactly at the line

A body at `L = 0` will not propose a change to its own arrangement, and this is not a failure of vision. It is the arithmetic of a group whose whole moment is consumed by its own maintenance. *Get out of your comfort zone* reverses the order of operations: the departure is dreamt from surplus, not from its absence.

Failure modes. *Demanding innovation from an `L = 0` body* — the correct move is to cut `C` first. *Reading `L = 0` as complacency* — it is capacity, measured. *Ordinal `M`* — demote per FVP-2.

See: Chapters 41, 47; Book One Node 00.

Class D — The lens at group scope

IF-15 . The group moment Statement.

```
M = (D / P) . E
```

`D` shared doctrine (the group's frame); `P` collective perception (its intake); `E` the environment; `M` the situation as the group constructs it. `D` and `P` must share units so that `r = D/P` is dimensionless; `M` inherits `E`'s units.

Derivation. CANON Addendum II §4, carried to group scope in Chapter 47. **This is the author's formula and it is a postulate, badged.** What is not postulated is the dimensional requirement, which is forced.

FVP.

Gate	Verdict
FVP-1	**PASS (conditional).** `D` and `P` in the same unit, `M` in `E`'s unit. **If `D` and `P` are in different units the expression is not defined**

Gate	Verdict
FVP-2	**DEMOTE** (Standing Verdict 1). Division requires `D` and `P` ratio-scaled with a true zero on a **common** scale (Stevens' typology), and no such instrument yet exists — channel-counting gives a zero within one encoder and no linkage between encoders. **This is the crux of the whole volume; Chapter 53 is about exactly this and returns DEMOTE.** The formula stands as the ordering *over-framed* / *calibrated* / *under-framed* plus its limits and fixed point; its magnitudes are withdrawn. **Route out: Ch 54 Thm 54.3**, Rasch logit difference given badged Postulate SO1, ~123 items per term for half a logit, with decisions as items
FVP-3	**PASS.** `P -> 0` gives `M -> infinity`: an institution that has stopped listening. `D -> 0` gives `M -> 0`: a body that measures everything and concludes nothing. `E -> 0` gives `M -> 0`. All three are real, nameable, and common
FVP-4	**PASS.** `D = P` gives `M = E`: accurate situational assessment, stated as an equation. **This is the formula's most important content**
FVP-5	**PASS (conditional).** `D` and `P` are separately identifiable only through the prediction-surprise ledger; from `M` alone only the ratio is recoverable, and `E` is confounded
FVP-6	**PASS (conditional).** Killed if groups at `r = 1` are not measurably better calibrated than groups at `r = 3` on held-out predictions
FVP-7	**OPEN**

Protocol. 1. **Prediction-surprise ledger.** Before each period, the group records its predictions with stated confidence. After, score them. 2. `D` is estimated from the confidence held; `P` from the information actually supplied by the period's observations. Chapter 47 gives the construction in full. 3. **Test for cyclicity first** (Chapter 50 §50.14 B3). If the group's collective ordering over the relevant options is intransitive, `D` is not a well-defined object and this entry does not apply.

Worked number. With `E = 10.0`:

D	P	r	M	reading
1.0	1.0	1.000	10.0000	calibrated: `M = E`
2.5	1.0	2.500	25.0000	over-framed
3.0	1.0	3.000	30.0000	over-framed, severe

D	P	r	M	reading
1.0	3.0	0.333	3.3333	under-framed, severe
0.4	1.0	0.400	4.0000	under-framed

Failure modes. *Ordinal inputs* — FVP-2; demote or do not publish. *Cyclic `D`* — the numerator does not exist. *Confusing `M` with `E`* — they are equal only at the fixed point, and the whole diagnostic is about the cases where they are not.

See: Chapters 34, 47, 53.

IF-16 . Doctrine-to-data and the group distortion Statement.

```
r     = D / P
Delta = |ln r|        the distortion, in nats

r > 1  OVER-FRAMED   doctrine outruns intelligence; the body sees what its model predicts
r = 1  CALIBRATED     contact
r < 1  UNDER-FRAMED   the body drowns in data it cannot read
```

Derivation. Definition from IF-15. **The log-ratio form is chosen because ratios compose multiplicatively and the whole framework is logarithmic** (CANON Addendum II §1); it makes over-framing by a factor of 2 and under-framing by a factor of 2 equal distortions, which the raw ratio does not.

FVP.

Gate	Verdict
FVP-1	**PASS.** `r` dimensionless; `Delta` in nats
FVP-2	**Inherits IF-15's FVP-2: DEMOTE** (Standing Verdict 1). **`Delta` is an ORDINAL quantity pending instrumentation — it ranks bodies and periods by distance from contact and does not license the claim that one distortion is twice another.** The **sign** of `ln r` is unaffected and it is the part that decides the correction. The band table below is a construction showing how the measure reads, not a set of measurements. Restored to ratio scale by Ch 54 Thm 54.3 and by nothing cheaper
FVP-3	**PASS.** `r -> 0` and `r -> infinity` both give `Delta -> infinity`, symmetrically, which is the point of the log form
FVP-4	**PASS.** `r = 1`, `Delta = 0`: contact

Gate	Verdict
FVP-5	**PASS (conditional).** `Delta` alone loses the sign; **always report `r` with `Delta`**, because the two faults have opposite corrections
FVP-6	**PASS.** Killed if over-framed and under-framed bodies show the same held-out prediction error profile
FVP-7	**OPEN**

Protocol. 1. Prediction-surprise ledger, per IF-15. 2. **Over-framing signature:** high stated confidence, flat accuracy — the body is as confident when wrong as when right. 3. **Under-framing signature:** low stated confidence with accuracy that would have supported more — the body had the information and did not use it. 4. Classify by band and report `r`, not only `Delta`.

Worked number. Bands:

`Delta = |ln r|`	over-framed `r`	under-framed `r`	reading
< 0.10	< 1.105	> 0.905	calibrated
0.10 – 0.40	1.105 – 1.492	0.905 – 0.670	drift
0.40 – 1.00	1.492 – 2.718	0.670 – 0.368	fault
> 1.00	> 2.718	< 0.368	severe

A body at `r = 2.90` has `Delta = 1.0647` — severe over-framing. Correcting to `r = 1.25` gives `Delta = 0.2231` — drift, which is the realistic target. **Calibration is not the goal; the achievable goal is drift, and a body reporting `Delta < 0.05` is more likely to be mismeasuring than to be perfect.**

Two structural reasons why `r = 1` is not an unqualified target, added here as corrections to any reading of this entry that treats the band table as an objective function.

(a) The fixed point converges to exactly zero headroom. At `r = 1` the input sits at its ceiling, `S = ln N`; Paper E Theorem 3.5 gives `Q_inf = (eta.ln N − C)/delta`, and IF-09's gauge then reads `h = eta.ln N − delta.Q_inf − C = 0` exactly, `delta` and `C` cancelling. **Perfect contact drives the input to its maximum, the maximum input drives the store to its maximum, and the maximum store consumes the whole budget in maintenance.** Chapter 47 §47.6 derives it; a body wanting headroom at equilibrium must shed `Q` deliberately or hold a declared `r != 1`. Note also the exact join to Paper A, `ln r = ln D − ln p_bar − F` (Chapter 47 §47.8): **the group's surplus enters the doctrine-to-data ratio one nat for one nat**, so calibration is a geometry requirement and is infeasible outright unless `D/p_bar <= N`.

(b) A body optimising raw `Q` is driven away from `r = 1`, actively. Over-framing **increases** stored order — the storage law has no term asking whether the moment was earned (Chapter 34 §34.21) — so the two-track ledger of Chapter 40 §40.2 is mandatory here: `Q` against veridical `Q^V`, stored error `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)`, correction load `C_t = c_0 + kappa.Z_t`. Since `V = E.min(r,1)` is flat above the fixed point while `M = rE` climbs, **dogma banks faster than knowledge and the bill arrives on a lag.** The knife edge is `kappa = delta`;

below it raw `Q_inf` rises without bound in `r`. **Report `Z` alongside `Q`, or report `Q^V`.** Chapter 47 §47.9.1 works this in full; Chapter 48 gives the corrections.

Failure modes. *Correcting the wrong direction* — under-framing treated by adding doctrine becomes over-framing; the sign matters more than the magnitude. *Reporting `Delta` alone* — loses the sign. *Confusing over-framing with confidence* — the diagnostic is confidence **uncorrelated with accuracy**, not confidence as such.

See: Chapters 35, 48.

IF-17 . The group return difference Statement.

```
Delta_R = | ln(M' / M) |
```

`M` the moment the group sent into the world (its action, its proposal, its signal); `M'` the moment that came back. Log-ratio, in nats. Convertible to octaves by `x 2.885390` (CANON Addendum II §1).

Derivation. Book One Node 08 ECHO, formalised in CANON Addendum II §8, carried to group scope. **The claim that the world's answer over what was sent is the informative quantity is a postulate, badged.** The log form is forced by the same reasoning as IF-16.

FVP.

Gate	Verdict
FVP-1	**PASS (conditional).** `M` and `M'` in the same unit — which is the hard requirement, because the answer often arrives in a different currency than the question
FVP-2	**DEMOTE** (Standing Verdict 1). The ratio requires `M` ratio-scaled and `M` inherits IF-15's demotion. **Amplification, null return and attenuation are findings; the number of nats is not.** The octave conversions below are exact arithmetic on an illustrative construction, not measurements. Ch 54 Thm 54.3 supplies the missing scale
FVP-3	**PASS.** `M' = M` gives `Delta_R = 0`: **no information came back.** `M' -> 0` or `M' -> infinity` gives divergence, corresponding to total rejection and to overwhelming response
FVP-4	**PASS.** `Delta_R = 0` is the fixed point and it is the *uninformative* one, which is the entry's central content
FVP-5	**PASS (conditional)** on `M` and `M'` being measured by the same instrument
FVP-6	**PASS.** Killed if `Delta_R` is uncorrelated with subsequent adjustment quality

Gate	Verdict
FVP-7	**OPEN**

Protocol. 1. Record `M` at the moment of sending, in the declared unit, **before** any response. 2. Record `M'` when the response is complete, same instrument. 3. Report magnitude **and sign of `ln(M'/M)`**: amplification versus attenuation are different findings.

Worked number. `M = 10`:

`M'`	`Delta_R` (nats)	octaves	reading
10	0.0000	0.0000	**zero return difference: no information**
12	0.1823	0.5261	mild amplification
7	0.3567	1.0291	attenuation
25	0.9163	2.6439	strong amplification
3	1.2040	3.4739	strong attenuation

A response identical to what was sent is an echo, and an echo carries no information about the world. A body whose consultations return exactly its own proposal has learned nothing, and the measurement says so with a zero.

Failure modes. *Treating `Delta_R = 0` as success* — the commonest misreading; it means the channel is a mirror. *Mismatched instruments* — measuring `M` by intent and `M'` by outcome makes the ratio meaningless. *Ignoring the sign* — amplification and attenuation call for opposite adjustments.

See: Chapters 43, 47; Book One Node 08.

Class E — Trust, reciprocity, obligation

IF-18 . Trust as delta reduction Statement.

```
delta_eff = delta_0 + (1 − a)(1 − v)
```

a the **reliance rate** — the measured probability that an *unverified* commitment is honoured within the period; v the fraction re-verified; `delta_0` the verification-irreducible loss.

Derivation. Chapter 49, Postulate T. **Badged: the mapping of trust onto `delta` is a modelling choice.** The decomposition of lapse into irreducible and confirmation-addressable causes is the substantive content, and the postulate is that verification is effective on the second class.

FVP.

Gate	Verdict
FVP-1	**PASS.** All three terms dimensionless per period

Gate	Verdict
FVP-2	**PASS.** a and v are proportions with true zeros — this is why trust was defined as a rate rather than as a belief
FVP-3	**PASS.** `a = 1` gives `delta_eff = delta_0` regardless of v; `v = 1` gives `delta_eff = delta_0` regardless of a. **The two controls substitute exactly, which is the entry's whole point**
FVP-4	**PASS.** `delta_eff = delta_0` is the floor, reached from either direction
FVP-5	**PASS (conditional).** a and v are separately identifiable only if some commitments are genuinely left unverified. **Under `v = 1`, a is unmeasurable**
FVP-6	**PASS.** Killed if agreements in high-a relations do not survive longer unmaintained than those in low-a relations
FVP-7	**OPEN**

Protocol. Chapter 49, Protocol A: 1. Enumerate the commitment population in one domain. **Failure at this step is itself the finding.** 2. Stratify by consequence. 3. Randomise non-verification **within the low-consequence tier only**. 4. `a_hat` = honoured / total among the unverified. Report a Wilson interval. 5. Repeat per domain; never pool.

Ethics, binding. Parties must be told that verification is sampled. **Which** commitments are unverified may be concealed; **that** sampling occurs may not. A hidden reliability test is a covert manipulation and is prohibited under CANON §5.4.

Worked number. `delta_0 = 0.04`, `v = 0`:

a	delta_eff	half-life of an unmaintained agreement
0.50	0.5400	0.893 periods
0.70	0.3400	1.668
0.80	0.2400	2.526
0.90	0.1400	4.596
0.95	0.0900	7.350
0.99	0.0500	13.513

Elasticity of `Q_inf` in a is `a/(1−a)`: 1.00 at `a = 0.5`, 19.00 at `a = 0.95`, 99.00 at `a = 0.99`. **The last increments of reliability are worth absurdly more than the first**, and the same arithmetic explains why one prominent breach in a high-a body is so destructive.

Failure modes. *Scalar trust* — a is indexed by pair and domain; the average predicts neither. *Trust from survey* — an ordinal self-report, prohibited by FVP-2 from entering the equation. *Maximising a* — `delta` is not a fault to be minimised (Chapter 3); high a on a commitment that ought to lapse is how obsolete procedure becomes immortal. Use sunset clauses.

See: Chapters 3, 49.

IF-19 . The verification budget and the trust threshold Statement.

```
C_v(v)  = c_v * v
Q_inf(v) = (K - c_v v) / (delta_0 + (1-a)(1-v)),     K = eta S - C_0

sign( dQ_inf/dv ) = sign[ (1-a)(K - c_v) - c_v * delta_0 ]    -- independent of v
a* = 1 - c_v * delta_0 / (K - c_v)        valid when K > c_v
```

Derivation. Chapter 49 §49.5. `Q_inf(v)` is linear-fractional in `v`, so its derivative has constant sign and **the optimum is always at an endpoint: verify everything or verify nothing.**

FVP.

Gate	Verdict
FVP-1	**PASS.** `c_v` in `Q`-units per period at full coverage
FVP-2	**PASS** given a ratio-scaled cost measure
FVP-3	**PASS.** `c_v >= K` means verification consumes the entire net input and is never worth it; `c_v -> 0` drives `a* -> 1`, verify almost always
FVP-4	**PASS.** `a = a*` is the indifference point
FVP-5	**PASS (conditional)** on independent estimates of `c_v`, `delta_0` and `K`
FVP-6	**PASS.** Killed if bodies verifying at the predicted-suboptimal rate show no `Q` penalty
FVP-7	**OPEN**

Protocol. 1. Estimate `c_v` as the full-coverage verification cost per period. 2. Estimate `delta_0` from an interval with `v = 1` (verification saturated, so residual loss is irreducible). 3. Estimate `K = eta S - C_0` per IF-06 and IF-02. 4. Compute `a*` and compare against measured `a` (IF-18).

Worked number. `K = 0.2526`, `delta_0 = 0.04`:

c_v	a*	policy
0.02	0.9966	verify unless the relation is essentially perfect
0.06	0.9875	verify below 98.75%
0.12	0.9638	verify below 96.4%
0.20	0.8479	verify below 84.8%
>= 0.2526	never	verification consumes the whole net input

Trust is valuable exactly in proportion to the cost of the verification it displaces. Where verification is cheap — a barcode scan — trust is a luxury and its absence costs nothing. Where verification is expensive or impossible — judgement, taste, discretion, intent — trust is the only available `delta` control and its absence is a hard ceiling on stored order.

Model artefact, flagged. The bang-bang result follows from the **linearity** of `C_v = c_v v`. Real verification costs are convex and the true optimum is interior. **Use the endpoint result as a signpost about direction, never as an instruction to audit everything or nothing.** What survives convexification is the threshold structure.

Failure modes. *Verification creep* — controls added after an incident and never sunset; price the current v against a∗. *Monitoring mistaken for trust-building* — monitoring lowers `c_v`, which **raises** a∗; it does not raise a (Ostrom's fourth principle, correctly read).

See: Chapter 49; Ostrom 1990.

IF-20 . The reciprocity index and the eta law Statement.

```
rho_ij = 2 min(x_ij, x_ji) / (x_ij + x_ji) = 1 − |x_ij − x_ji| / (x_ij + x_ji)
rho_group = ( sum over pairs of matched volume ) / ( total directed volume )
eta = eta_max ∗ rho
```

The two expressions for `rho_ij` are identical, since `a + b − |a−b| = 2 min(a,b)`. **`rho` is simultaneously the matched fraction of the volume and one minus the normalised imbalance.**

Derivation. The index is a definition. **`eta = eta_max ∗ rho` is Chapter 49's Postulate R, badged**: only the matched fraction converts, because structure only one party holds is a debt rather than a structure.

FVP.

Gate	Verdict
FVP-1	**PASS.** `rho` dimensionless; invariant to the flow unit — required before it can multiply `eta_max`
FVP-2	**PASS (conditional).** Requires a **symmetric** ratio-scaled flow unit: the same unit in both directions. **If you count advice one way and money the other, `rho` is undefined**
FVP-3	**PASS.** `x_ji = 0` gives `rho = 0` and `eta = 0`: a one-way relation deposits nothing durable. `x_ij = x_ji` gives `rho = 1`
FVP-4	**PASS.** `rho = 1` is the balanced fixed point
FVP-5	**FAIL as stated: `rho` depends on the accounting window.** The same relation measures 0.13 at a one-month window and 0.87 at twenty-four months. **The window must be preregistered and the sensitivity curve reported**
FVP-6	**PASS.** Killed if the shared structures that survive a full year are uncorrelated with `rho`
FVP-7	**OPEN**

Protocol. Chapter 49, Protocol B: 1. Declare the symmetric ratio-scaled flow unit. 2. **Declare the**

window before looking. Default `w = 1/delta` periods: beyond the time constant, the structure a return was meant to build has already leaked away. 3. Build the full directed matrix `x_ij`. 4. Compute `rho` per dyad and for the group. 5. Report the window sensitivity curve and the network density alongside. **A network with three edges can have `rho = 1`.**

Worked number. `S = 1.5790`, `C = 0.30`, `delta = 0.10`, `eta_max = 0.45`:

x	y	rho	eta	eta S	Q_inf
10	10	1.0000	0.4500	0.7106	+4.1055
9	6	0.8000	0.3600	0.5684	+2.6844
8	3	0.5455	0.2455	0.3876	+0.8757
12	2	0.2857	0.1286	0.2030	**-0.9699**
5	0	0.0000	0.0000	0.0000	**-3.0000**

A relation at `rho = 0.29` does not accumulate slowly. It fails to accumulate at all — and it does so while fourteen transfers a period look perfectly healthy on a dashboard.

Window sensitivity, fifteen favours with mixed return lags:

window (months)	1	2	3	6	12	24
estimated `rho`	0.1333	0.2667	0.4000	0.6000	0.8000	0.8667

Failure modes. *Window shopping* — a `rho` that improved after the analyst "refined the measurement period"; ask when the window was chosen. *Asymmetric units* — FVP-2. *Volume mistaken for reciprocity* — the dashboard measures the total and the store depends on the matched fraction. *Ignoring density* — `rho` on a sparse network is a statement about the few edges that exist.

See: Chapters 4, 46, 49.

IF-21 . The obligation decomposition and the imbalance index Statement.

```
X = Sym + Skew,   Sym_ij = (x_ij + x_ji)/2,   Skew_ij = (x_ij - x_ji)/2
b_i = sum_j (x_ij - x_ji)        node net obligation balance, sum_i b_i = 0
B   = sum_i |b_i| / (2 * total volume)
theorem:  1 - B  >=  rho
```

Derivation. The decomposition is exact and unique for any square matrix. **The interpretation — `Sym` measures the relation and `Skew` measures the obligation — is Chapter 49's Postulate O, badged.** The bound is proved: `|sum_j (x_ij - x_ji)| <= sum_j |x_ij - x_ji|`, so node-level imbalance never exceeds dyad-level imbalance, hence `1 - B >= rho`.

FVP.

Gate	Verdict
FVP-1	**PASS.** `Skew` in flow units; `B` dimensionless
FVP-2	**PASS (conditional).** Same symmetric ratio-scale requirement as IF-20
FVP-3	**PASS.** `B = 0` is node-balance; `B -> 1` is pure one-directional flow
FVP-4	**PASS.** `Skew = 0` is the reciprocal fixed point, where the relation is entirely `Sym`
FVP-5	**PASS (conditional).** `B` is identifiable; **`rho` is not recoverable from `B`** — only bounded, which is the theorem
FVP-6	**PASS.** Killed if dyads ranked by `|Skew|` do not predict breakdown better than dyads ranked by `Sym`
FVP-7	**OPEN**

Protocol. Chapter 49, Protocol C: 1. Same matrix as IF-20 — one collection serves both. 2. Compute `Skew`, `b_i`, `B`. 3. **Rank dyads by `|Skew|`, not by `Sym`.** The largest relationships are not the most imbalanced. 4. **Check the sign structure.** A few large net creditors and many small net debtors is a patron-client structure whatever it calls itself; `b_i` near zero with high `Sym` is a horizontal association. These are Putnam's (1993) two cases, distinguishable in one afternoon of arithmetic.

Ethics, binding. Individual balances shall not be published within the group. A public ledger of who owes whom converts a private accounting into a public pressure and removes the ability to decline the accounting. Publish the distribution; report dyad findings to the dyad.

Worked number. Eight-firm consortium, total directed volume 68:

firm	F1	F2	F3	F4	F5	F6	F7	F8
`b_i`	-17	-21	+14	+8	+5	+5	+3	+3

```
B     = 76 / (2 * 68) = 0.5588
1 - B = 0.4412        measured rho = 0.3824
```

The node index flatters the network by fifteen per cent, and it does so systematically, in the direction of comfort. Report `rho`, or report both.

Failure modes. *Reporting `B` alone* — biased toward good news. *Publishing the ledger* — see the ethics clause. *Reading a large `Sym` as health* — volume and balance are independent.

See: Chapters 46, 49; Gambetta 1993; Putnam 1993.

Class F — Members and thresholds

IF-22 . Member headroom and the exit condition Statement.

```
h_i = eta ln N / N  −  delta Q / N  −  c_i
critical private cost:  c_i∗ = (eta ln N − delta Q)/N
member i is individually unsustainable when h_i < 0, regardless of the group's h
```

Derivation. IF-09 divided by N, with c_i the member's own private cost share replacing the per-capita average. **The equal division of the ceiling and of maintenance is a postulate, badged**; unequal division is the general case and requires member-specific shares.

FVP.

Gate	Verdict
FVP-1	**PASS.** Q-units per period per member
FVP-2	**PASS** given IF-06's condition and a ratio-scaled c_i
FVP-3	**PASS.** c_i = 0 gives the per-capita headroom; c_i large gives h_i arbitrarily negative
FVP-4	**PASS.** c_i = c_i∗ is the individual break-even
FVP-5	**PASS (conditional)** on measuring c_i per member rather than assuming the average
FVP-6	**PASS.** Predicts which members leave and in what order. **Directly and cheaply falsifiable**
FVP-7	**OPEN**

Protocol. 1. Measure c_i per member: their own time, travel, emotional labour, compliance burden, in the group's Q-unit. 2. Compute h_i for each. 3. **Report the names below zero, not just the count.** 4. Address by reducing c_i — never by restricting exit. **CANON §5.4 governs: retention is an offer, never a restriction.**

Worked number. N = 9, u = 0.6, eta = 0.35, delta = 0.08, Q = 4.0. Per-member ceiling share 0.08545; per-member maintenance share 0.03556; so c_i∗ = 0.04989:

member	c_i	h_i	status
1	0.02	+0.02989	ok
2	0.03	+0.01989	ok
3	0.05	-0.00011	unsustainable
4	0.09	-0.04011	unsustainable
5	0.14	-0.09011	unsustainable
6	0.22	-0.17011	unsustainable

Four of six below water in a body that may be reporting positive aggregate headroom. They leave in cost order — highest c_i first — and the highest-cost members are frequently the highest-u members, because carrying a divergent framing in a converged room is expensive. **Combine**

with IF-05: the body loses its most divergent members first and therefore loses surplus faster than headcount.

Failure modes. *Aggregate headroom quoted as sufficient* — it is not. *Attributing exit to commitment* — `h_i < 0` is arithmetic. *Retention by restriction* — prohibited.

See: Chapters 6, 22; Chapter 50 §50.14 B1; Hirschman 1970.

IF-23 . Rival versus non-rival return Statement.

```
RIVAL     (store divided):    return_i = eta ln(N_eff) / N  -  c_i
NON-RIVAL (public good):      return_i = eta ln(N_eff)      -  c_i

ln(x)/x is maximised at x = e = 2.71828
```

Derivation. Definition of the two allocation regimes, applied to the surplus of IF-02. **The maximum of `ln(x)/x` at `x = e` is elementary calculus and it is exact.**

FVP.

Gate	Verdict
FVP-1	**PASS**
FVP-2	**PASS** given a ratio-scaled return measure
FVP-3	**PASS.** Rival: `N -> infinity` gives `return -> -c_i`. Non-rival: `N -> infinity` gives `return -> infinity`
FVP-4	**PASS.** Rival optimum at `N_eff = e`; non-rival has no interior optimum
FVP-5	**PASS (conditional).** Rivalry must be **declared** from the nature of the store, not inferred from the data — both regimes fit small-N data equally well
FVP-6	**PASS.** The two regimes make opposite predictions about the size gradient. **This is the cleanest falsification test in the chapter**
FVP-7	**OPEN**

Protocol. Declare rivalry first (Chapter 50, S50.5): is the store divided among members or held whole? Tooling access is rival; a published standard is not. Measure `c_i` and the outside option `w_i` per member, predict the size gradient, and check it against observed joining and leaving.

Worked number. `eta = 0.35`, `c_i = 0.020`, outside option `w = 0.030`.

Rival:

N	3	6	12	20	30
per-capita	0.12817	0.10452	0.07248	0.05243	0.03968

N	3	6	12	20	30
return	+0.10817	+0.08452	+0.05248	+0.03243	+0.01968
decision	stay	stay	stay	stay	**exit**

Non-rival:

N	3	5	8	12	20	30
return	+0.3645	+0.5433	+0.7078	+0.8497	+1.0285	+1.1704

The same group, the same surplus, the same members — and the sign of the size gradient flips entirely depending on whether the store is divided or shared. Rival stored order pushes groups toward about three; non-rival stored order has no upper limit from this term and is bounded only by `C` (IF-12).

Failure modes. *Undeclared rivalry* — predictions about optimal size are unsupported without it. *Treating a mixed store as one regime* — most real institutions hold both, and the correct move is to split `Q` into two accounts. *Reading the `x = e` result as a recommendation* — it is the rival regime's arithmetic, not a design rule.

See: Chapters 5, 22; Chapter 50 §50.14 B8; Ostrom 1990.

IF-24 . The tipping threshold and the cascade Statement.

`committed minority required to overturn an established convention: ~25% of N`
`Granovetter threshold cascade: with thresholds T_i, the equilibrium active count x* is the smallest fixed point of  x -> #{ i : T_i <= x }`

Derivation. Centola, Becker, Brackbill & Baronchelli (*Science* 360, 2018) **measured** a tipping point near 25 per cent for the overturning of an established social convention. Granovetter (*AJS*, 1978) gives the threshold-model fixed-point construction. **Both are real results; the transfer to your body is a modelling assumption, badged.**

FVP.

Gate	Verdict
FVP-1	**PASS.** Counts and proportions
FVP-2	**PASS.** Proportion of a countable membership
FVP-3	**PASS.** All thresholds zero gives full cascade immediately; all thresholds >= N gives none
FVP-4	**PASS.** The equilibrium is a fixed point of the activation map — **the same fixed-point structure as the storage law, arrived at from a different direction**

Gate	Verdict
FVP-5	**PASS (conditional).** Thresholds are hard to measure directly; they are usually inferred from observed cascades, which is circular unless the inference is out of sample
FVP-6	**PASS (conditional).** The 25 per cent figure is one well-conducted experimental paradigm, not a constant of nature. **State it as a measured tipping point in that paradigm**
FVP-7	**PARTIAL.** Centola's result is a published controlled experiment with a clear design; broad cross-context replication of the specific 25 per cent value is **not** established and must not be claimed

Protocol. 1. Count `N` for the convention in question — the population the convention actually binds. 2. Committed count required: `ceil(0.25 N)`, treated as an order-of-magnitude guide, not a target. 3. **Estimate the threshold distribution** by observing partial adoptions, and compute the fixed point. The distribution matters more than the mean.

Worked number.

N	8	12	20	25	40	100	150
committed needed at 25%	2	3	5	7	10	25	38

Granovetter's point, computed. Ten people with thresholds `0,1,2,3,4,5,6,7,8,9` cascade to **10 of 10**. Change one threshold — the person at 1 becomes a 2 — and the equilibrium is **1 of 10**.

One person's threshold moving by one unit changes the outcome from complete adoption to none. The mean threshold moved by 0.1; the outcome moved by 90 per cent of the population. This is why aggregate measures of "readiness for change" have almost no predictive value and why the distribution is the only useful object.

Complex contagion caveat, mandatory. Centola & Macy (*AJS*, 2007) and Centola (*Science*, 2010) established that behaviours requiring social reinforcement spread differently from simple contagion: **they need multiple reinforcing contacts, and clustered topology beats random long ties.** A change requiring reinforcement will not propagate along the weak-tie shortcuts that carry information, and a plan built on information diffusion will fail for that reason without the failure being visible.

Failure modes. *Quoting 25 per cent as a law* — FVP-6 and FVP-7. *Manufacturing the committed minority* — a committed minority assembled by instruction is conscripted difference (IF-01) and, where members cannot decline the role, is prohibited by CANON §5.4. *Planning simple diffusion for a reinforcement-dependent change* — see the caveat.

See: Chapters 15, 19; Centola et al. 2018; Granovetter 1978; Centola & Macy 2007.

IF-25 . The scope transform Statement.

```
delta_T = 1 – (1 – delta)^T       aggregating T short periods into one long period
delta   = 1 – (1 – delta_T)^(1/T)   the inverse
```

Period-free quantities: N, u, N_eff, F. Period-bearing: delta, eta, C, h, and Q if it is a flow rather than a stock.

Derivation. Survival compounds multiplicatively: the probability of surviving T independent periods is (1–delta)^T. **Exact, not an approximation.**

FVP.

Gate	Verdict
FVP-1	**PASS.** Both sides dimensionless
FVP-2	**PASS.** Proportions
FVP-3	**PASS.** T = 1 is the identity; T –> infinity gives delta_T –> 1 for any delta > 0
FVP-4	**PASS.** delta = 0 is the only fixed point across all T
FVP-5	**PASS.** Directly computable
FVP-6	**PASS.** Killed if measured multi-period decay departs systematically from the compounded single-period rate — which **would** indicate heterogeneity (Chapter 50 §50.14 B2) rather than a failure of the arithmetic
FVP-7	**OPEN**

Protocol. Convert **before** comparing any two rates. Never compare delta values whose periods differ.

Worked number.

delta per period	T = 7	T = 30	T = 52	T = 365
0.001	0.0070	0.0296	0.0507	0.3059
0.005	0.0345	0.1396	0.2295	0.8395
0.010	0.0679	0.2603	0.4070	0.9745
0.020	0.1319	0.4545	0.6503	0.9994
0.050	0.3017	0.7854	0.9306	1.0000
0.100	0.5217	0.9576	0.9958	1.0000

The trap, computed:

```
internal delta = 0.02 per DAY,          compounded over 91 days   –
> 0.8409 per quarter
external delta = 0.02 per QUARTER, expressed per day      –> 0.000222 per day
ratio of implied daily rates                                –> 90.1x
```

Two numbers that print identically differ by a factor of ninety.

Failure modes. *Unnormalised comparison* — the single commonest cross-scope error. *Adding rates instead of compounding* — at `delta = 0.10` and `T = 7`, addition gives 0.70 and the truth is 0.5217. *Applying the transform to a heterogeneous store* — the compounding assumes one `delta`; if the fit departs from the prediction, that is the diagnosis, not a rounding problem.

See: Chapters 2, 3; Chapter 50 §50.14 B7.

The Specification — using and extending the formulary

S51.1 — FVP row mandatory. No formula shall be used in a report from this book without its FVP row reproduced or cited. **An entry quoted without its gates is an entry quoted without its limits.**

S51.2 — OPEN is a debt, not a pass. Twenty-four of the twenty-five entries have FVP-7 OPEN and the twenty-fifth has it only PARTIAL. Any claim of established measurement shall state that. Chapter 56 sets out what closing that gate would require.

S51.3 — Demote on scale failure. Where FVP-2 fails — ordinal data entering a division or a ratio — the formula shall be demoted in print to a heuristic ordering, labelled as such, and its numerical output shall not be published. This governs IF-14, IF-15, IF-16 and IF-17 most often.

S51.4 — Badge every postulate in the sentence that states it. IF-04, IF-05, IF-11, IF-13, IF-14, IF-15, IF-17, IF-18, IF-20, IF-21, IF-22, IF-23 and IF-24 all rest on at least one postulate. **The mysticism is permitted, but it is permitted in the open, with its badge on**, and so is the modelling.

S51.5 — Report the ceiling and the realised fraction together. Any use of IF-09 shall be accompanied by `S / ln N` from IF-02. **Half of all misdiagnosis turns on confusing the ceiling with the realised value.**

S51.6 — Preregister windows and partitions. IF-03's block partition and IF-20's accounting window shall be declared before the data are seen, with justification.

S51.7 — Distribution, not scalar. IF-05, IF-11, IF-20 and IF-22 shall be reported with their across-member distributions. A scalar alone is a mean-field claim.

S51.8 — Consent constraints carried. IF-18's disclosure requirement, IF-21's ledger-privacy requirement, IF-22's no-restriction requirement, and IF-24's no-conscription requirement travel with those entries and shall not be dropped in citation.

S51.9 — No chaining across triples. IF-03 shall not be composed across three parties. The triangle inequality fails (Paper A §5).

S51.10 — Adding an entry. A new formula enters this formulary only with: a stated derivation identifying it as theorem, definition or postulate; all seven FVP gates with verdicts; a measurement protocol executable by someone who did not write it; at least one computed worked number; and named failure modes. **A formula without a failure mode has not been thought about.**

The Failure Modes — of formulary use

FM-51.1 — The Quoted Number. *Signature:* a figure from a worked example in this chapter reappearing in someone's report as a finding about their system. *Damages:* every downstream conclusion. *Diagnostic:* ask which of their own measurements produced it. Usually none. *Fix:* every entry's worked number is an illustration of the arithmetic, not a benchmark. Measure.

FM-51.2 — Gate Skipping. *Signature:* an entry applied with its FVP row omitted, most often FVP-2. *Damages:* produces confident numbers from ordinal data. This is the single fault that separates a measured field from an unmeasured one (Chapter 52). *Diagnostic:* ask for the scale type of every term. If the answer is "a 1–5 rating," the division is invalid. *Fix:* S51.1 and S51.3.

FM-51.3 — Postulate Laundering. *Signature:* a badged postulate quoted at second hand without its badge, and then treated as following from Paper A or Paper E. *Damages:* misattributes a modelling choice to a theorem, which makes it unarguable. *Diagnostic:* trace the claim to its entry and check the Derivation field. *Fix:* S51.4. The four things this book refuses to do (CANON §5) exist for this.

FM-51.4 — The Single Lever. *Signature:* a diagnosis naming one parameter, and an intervention on that parameter alone. *Damages:* Chapter 49 §49.14 computed a case where a single-lever trust intervention **moved the equilibrium the wrong way** while every survey measure improved. *Diagnostic:* compute the one-at-a-time decomposition. If the interaction term exceeds either main effect, single-lever intervention is contraindicated. *Fix:* check the sign of `eta S − C` before any `delta` intervention (IF-08, IF-18, IF-19).

FM-51.5 — Ceiling/Realised Confusion. *Signature:* a body reported as having ample headroom and shrinking every period. *Damages:* sends the engineer to the wrong lever entirely. *Diagnostic:* h uses `ln N`; the growth condition uses `ln N_eff`. Compute both, always. *Fix:* S51.5.

FM-51.6 — Scalar Reporting. *Signature:* one u, one `delta`, one `rho` for a group with visibly heterogeneous members. *Damages:* conceals the fragility profile (IF-05) and the exit profile (IF-22), both of which are computable in an afternoon. *Diagnostic:* recompute with each member deleted in turn. If the spread is wide, the scalar was hiding it. *Fix:* S51.7.

FM-51.7 — Consent Stripping. *Signature:* IF-18 used to run hidden reliability tests; IF-21 used to publish a debt ledger; IF-22 used to justify retention clauses; IF-24 used to assemble a committed minority by instruction. *Damages:* violates CANON §5.4, and in three of the four cases also invalidates the measurement. *Diagnostic:* ask whether the affected person could have declined and been told they could. *Fix:* S51.8. **This is the one failure mode with no parameter fix and no acceptable trade-off.**

FM-51.8 — Reference Chapter Read As Argument. *Signature:* this chapter cited as though the density of formulas were itself evidence. *Damages:* the credibility of everything else. *Diagnostic:* count the OPEN gates. Of twenty-five FVP-7s, twenty-four are OPEN and the twenty-fifth is only PARTIAL. *Fix:* S51.2, and Chapters 52 through 56, which are about exactly this and are the honest answer to it. **Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does.**

Worked Example — A Regional Volunteer Rescue Association, Diagnosed Through Eleven Entries

A volunteer search-and-rescue association: **fourteen members on the roster**, across a county, nine years old, self-governing, donation-funded. Membership is voluntary in the fullest sense and members may leave at any time. The complaint: "we are losing our best people and we cannot work out why, because morale is good."

Period: one quarter. `Q` unit: retained-practice units. Rivalry declared (IF-23): **non-rival** — the store is trained capability and shared protocol, held whole by every member.

Step 1 — IF-02, the effective count

Capability inventory: **`N = 9`** domains verified by observed independent action (rope, water, medical, canine, navigation, communications, incident command, vehicle, public liaison) — five of the fourteen members duplicated domains already covered. Pre-discussion prediction round on eight past incidents gave **`u = 0.42`**.

```
N_eff = 1 + 8 * 0.42 = 4.360        S = ln 4.360 = 1.4725
ln N  = ln 9 = 2.1972               S / ln N = 67.0%
```

Fourteen people, nine domains, an effective count of 4.36. **Not the fault, but the number that every subsequent calculation depends on.**

Step 2 — IF-11, turnover composition

Four departures over four quarters against an average roster of 14 gives `tau = 0.0714` per quarter. Intrinsic decay per IF-07 over a two-quarter stable interval: `delta_int = 0.06`.

```
delta_total = 1 − (1 − 0.06)(1 − 0.0714) = 0.1271
```

Per IF-11's table structure, turnover is `0.0714/0.1271 = 56 per cent of total loss`. **More than half the association's forgetting is people leaving.**

Step 3 — IF-06 and IF-08, the storage law and the steady state

`eta` per Chapter 4 over eleven quarters: **`eta = 0.31`**. `C` as coordination plus compliance plus travel: **`C = 0.42`**.

```
eta*S  = 0.31 * 1.4725 = 0.4565
Q_inf  = (0.4565 − 0.42)/0.1271 = +0.2868
```

Positive, but barely. Current `Q = 1.85`, well above the equilibrium — **the association is declining toward 0.2868 and will take `ln2/−ln(1−0.1271) = 5.10` quarters to cover half the distance.**

Step 4 — IF-09, headroom

```
eta ln N = 0.31 * 2.1972 = 0.6811
h at Q = 1.85 = 0.6811 − 0.1271*1.85 − 0.42 = 0.6811 − 0.2352 − 0.4200 = +0.0259
```

Headroom of +0.0259 — effectively maintenance only. Per IF-09's regime table this is *hold, do not build*, and the association had spent the year planning an expansion.

Ossification point: `Q = (0.6811 − 0.42)/0.1271 = 2.0539`. **The association's Q of 1.85 is 90 per cent of the way to its own ceiling**, which is IF-09's structural result: success consumes its own headroom.

Step 5 — IF-10, the requested increment

The expansion plan requested `r = 0.15` Q-units per quarter.

```
required S = (0.1271∗1.85 + 0.42 + 0.15)/0.31 = (0.2352 + 0.42 + 0.15)/0.31 = 2.5975
ln N = 2.1972  ->  required S exceeds the ceiling.  INFEASIBLE.
required u = (e^2.5975 − 1)/8 = 1.5537 > 1.  Confirmed.
r_max = h = 0.0259
```

The plan asked for 0.15 and the maximum available is 0.0259 — a factor of 5.79. Not closeable by effort. Per IF-10, the only remaining moves are to raise N or lower C.

Step 6 — IF-22, member headroom

Ceiling share `0.6811/14 = 0.04865`; maintenance share `0.2352/14 = 0.01680`; critical private cost `c_i∗ = 0.03185`. Measured `c_i` (travel, training, callout burden, in Q-units):

member group	count	c_i	h_i	status
town-based, single domain	6	0.014	+0.01785	ok
town-based, multi-domain	3	0.028	+0.00385	marginal
rural, multi-domain	3	0.052	**-0.02015**	**unsustainable**
rural, incident command	2	0.071	**-0.03915**	**unsustainable**

Five of fourteen are individually below water while aggregate h is positive, and the four who left in the last year were all from the two rural groups. **The association was not losing people to a morale problem. It was losing them to a travel and callout burden that made their private headroom negative.** Morale read as good because the town-based majority — nine of fourteen — was above water.

Step 7 — IF-05, departure fragility

Per S51.7, u recomputed with each member deleted:

removed	resulting u	S at N unchanged	loss
any town-based single-domain member	0.42	1.4725	0.0000
a town-based multi-domain member	0.40	1.4351	0.0374
a rural multi-domain member	0.35	1.3350	**0.1375**

removed	resulting u	S at N unchanged	loss
a rural incident-command member	0.31	1.2470	**0.2254**

The members with the most negative `h_i` are precisely the members whose departure costs the most surplus. IF-05 and IF-22 compose, and that composition is the association's actual failure mode: **a private-cost gradient selecting for the departure of the highest-u members.** Random departure costs on average 0.0697; targeted departure of an incident-command member costs 0.2254 — **3.24 times as much**. The association had been experiencing targeted departure with nobody targeting anything.

Step 8 — IF-20 and IF-21, reciprocity and obligation

Flow matrix over four quarters, unit = **callout-hours covered for another member**, symmetric and ratio-scaled. Per S51.6 the window was preregistered at `w = 1/delta = 7.9` quarters and truncated to the four quarters of record; the truncation is reported.

```
total directed volume = 214 hours
matched volume        = 116 hours
rho_group = 116/214 = 0.5421
```

Node balances: the two rural incident-command members are net creditors at +31 and +24 hours; the six town-based single-domain members are net debtors between −4 and −11.

```
B     = 118/(2*214) = 0.2757        1 - B = 0.7243
```

Per IF-21's theorem, `1 - B = 0.7243 >= rho = 0.5421`: **the node index overstates by 34 per cent**, and the sign structure is IF-21's patron-client pattern — a few large net creditors carrying many small net debtors. **The people leaving were the creditors.**

Back-solving IF-20's eta law from the measured `eta = 0.31` gives `eta_max = 0.5719`. **At `rho = 0.85`, `eta` would be 0.4861** — 57 per cent more conversion with no change in headcount, domains, or u.

Step 9 — IF-16, doctrine-to-data

Prediction-surprise ledger over eleven incidents: confidence high and uncorrelated with accuracy — the over-framing signature. Cyclicity tested first per IF-15: the association's ordering over deployment options was transitive, so D was well defined and the diagnostic applied.

```
D = 1.85,  P = 1.00,  r = 1.850,  Delta = |ln 1.850| = 0.6152    FAULT band
```

Step 10 — the intervention, priced

change	entry	parameter	effect
Callout rota rebalanced so town members cover rural incidents in proportion	IF-20	rho 0.5421 -> 0.82	eta 0.31 -> 0.4690
Travel and equipment reimbursement for rural members	IF-22	c_i -0.021 for the five below water	all h_i positive
Two recurring meetings deleted; incident debrief kept	IF-12	C 0.42 -> 0.34	
Two duplicated town members cross-trained into uncovered domains	IF-02	N 9 -> 11	ln N 2.1972 -> 2.3979
Expansion plan withdrawn	IF-10	—	target was infeasible by 5.8x

Recomputed, with u held at 0.42 (deliberately not assumed to improve):

```
S      = ln(1 + 10*0.42) = 1.6487        S/ln N = 68.8%
eta*S  = 0.4690 * 1.6487 = 0.7732
delta  = assume tau falls to 0.02 with the five retained:
         delta_total = 1 - (0.94)(0.98) = 0.0788
Q_inf  = (0.7732 - 0.34)/0.0788 = +5.4977
h at Q = 1.85: 0.4690*2.3979 - 0.0788*1.85 - 0.34 = 1.1246 - 0.1458 -
0.3400 = +0.6388
r_max  = 0.6388,  against the withdrawn request of 0.15
```

Steady state from +0.2868 to +5.4977. Headroom from +0.0259 to +0.6388 — the withdrawn expansion is now feasible four times over. And it became feasible by paying five people's travel costs and deleting two meetings.

Step 11 — falsification, declared in advance

1. If rho_group does not exceed 0.70 within three quarters, the rota change failed and must be reported as unsuccessful.
2. If tau does not fall below 0.04 within four quarters, the h_i diagnosis (IF-22) is wrong in this association and the departures have another cause.
3. If Q does not rise above 2.60 within 1/delta = 12.7 quarters, the fitted parameters do not describe this system.
4. **Shuffled-input control, mandatory.** Re-fit with the quarterly U_t sequence permuted. If permuted U predicts Q as well as the real sequence, the causal story is a trend.
5. If a rural incident-command member departs and u does not fall by approximately 0.11, IF-05's fragility computation is wrong here.

The rule, extracted

Eleven entries, one afternoon of arithmetic, and a diagnosis that inverted the presenting complaint. The association believed it had a morale problem in a healthy body. It had a **private-cost gradient in a body at 90 per cent of its ossification point**, selecting for the departure of exactly the members whose departure cost the most surplus, in a network whose reciprocity index was 0.54 while its node-level index read 0.72. **None of those facts is visible without the formulary, and every one is computable from records the association already kept.**

A formulary is not a set of answers. It is a set of questions with their arithmetic attached, and the arithmetic is what stops you from believing the first story the system tells you about itself.

Part X — Validation

Chapter 52 . The Formula Validation Protocol

The Law

This chapter and the two that follow it are the scientific warrant of Volume Two. Everything before them is a construction. Everything after them depends on whether the construction is allowed to stand. If Part X fails, Parts VI through IX are an elaborate notation for opinions, and the honest thing to do would be to say so and stop.

So it is worth being exact about what is being claimed, and about what the objection actually is.

52.0.1 The objection, stated fairly

The objection is not that psychology and sociology are pseudoscience. **That charge is false, it is lazy, and this manual does not make it.** It is worth spelling out why, because a great deal of sloppy work — including sloppy work that flatters this manual's own project — hides behind the sloppiness of the charge.

Psychology and sociology are sciences. They have theories that make predictions, experimental methods, measurement instruments, institutions of criticism, and findings that hold. The contact hypothesis (Allport 1954; Pettigrew & Tropp's meta-analysis of over five hundred studies) is a real finding. Woolley and colleagues' collective intelligence factor (*Science*, 2010) is a real finding. Centola's tipping point at roughly twenty-five percent (*Science*, 2018) is a measured quantity from a controlled experiment. Ostrom's design principles were extracted from field data across decades and earned a Nobel Prize. These are not vibes.

What these fields have is **a measurement problem and a replication problem**, and those are different diseases from fraudulence, with different treatments.

52.0.2 The replication numbers, stated exactly

The Open Science Collaboration's *Estimating the reproducibility of psychological science* (*Science*, 2015) attempted direct replications of one hundred studies published in three high-ranking psychology journals. The headline result, as recorded in this manual's canon:

```
originals reporting a statistically significant result:      97 of 100
replications reporting a statistically significant result:   about 36 of 100
replication effect sizes:                                    roughly half the originals
```

Three pieces of arithmetic on those figures, computed:

```
36 / 97  = 0.3711     of the originally significant findings, about 37% replicated
97 - 36   = 61           percentage-point gap between claimed and recovered
halving r: r -> r/2  =>  r^2 -> r^2 / 4      variance explained falls to 25%
```

That last line is the one people miss. **Effect sizes "roughly halved" means variance explained roughly quartered**, because the variance a correlation accounts for goes as its square. An effect that explained 16% of variance in the original explains 4% in the replication, whatever the p-value says. Verified across the plausible range:

original r	original r^2	halved r	halved r^2	ratio
0.20	0.0400	0.100	0.0100	0.250
0.30	0.0900	0.150	0.0225	0.250
0.40	0.1600	0.200	0.0400	0.250
0.50	0.2500	0.250	0.0625	0.250

52.0.3 What the numbers are evidence of, and what they are not

They are **not** evidence that psychology is fake. Some failed replications are themselves false negatives, some are context-sensitivity rather than error, and one replication is not a verdict. They are also not evidence about sociology, which was not the sample.

They are evidence of something narrower and more actionable. **A field in which 97% of published results are significant and 36% of them survive a direct redo is a field whose publication filter is selecting on something other than truth.** That is a process fault, and process faults have process fixes: preregistration, adequate power, held-out prediction, and — the concern of these three chapters — better measurement.

Here is the observation that makes the whole of Part X possible, and it is the load-bearing observation of the chapter. **Psychophysics does not have this problem.** Weber's ratio law has been reproduced continuously since 1834. Fechner's logarithmic relation since 1860. Stevens' magnitude-estimation exponents replicate across laboratories, across decades, and across languages, and where they vary they vary in known and documented ways with method, range and instruction (CANON_INTERNAL §3). Psychophysics studies *subjective experience* — the loudness of a tone, the heaviness of a weight, the brightness of a light — which is exactly the class of quantity the sceptic says cannot be measured. And it replicates.

So the discriminating variable is not subjectivity. Psychophysics is as subjective as anything in Volume Two. The discriminating variable is **how the quantity is put on a scale.** Psychophysics built ratio scales, with a true zero, using procedures that recover the same exponent from different observers on different days in different rooms. Much of the rest of the field uses seven-point ratings, treats them as though they were metres, and averages them.

52.0.4 The claim of Part X, stated so it can be checked

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does.

That sentence is the thesis of Chapters 52, 53 and 54, and it cuts against this manual as much as for it. Writing `M = (D/P)E` does not make the study of moments a science. Nothing about the typographic dignity of an equals sign confers empirical status. An equation is a claim about relations among *measured quantities*, and if the quantities are not measured — really measured, on a scale that supports the operations the equation performs on them — then the equation is a diagram of an intuition wearing mathematical clothes.

This is achievable. Psychophysics achieved it for sensory magnitude, and Rasch measurement achieved a version of it for attainment and attitude (Chapter 54). The Formula Validation Protocol is the route, and it is deliberately harsh: a protocol this manual's own formulas cannot fail would be worthless as a certificate that they passed.

52.0.5 The protocol, badged

Postulate FVP (methodological). **A formula in this manual is admissible if and only if it passes all seven gates below, in print, with its results stated.** This is a house standard, not a theorem. It is not derived from Paper A or Paper E, and no result in the corpus depends on it. It is a rule the author has adopted and by which the author asks to be judged. Its components are drawn from established measurement theory and established replication practice; the choice to require all seven, together, of every formula, is the author's.

```
FVP-1  DIMENSIONS      Units of every term; both sides agree; ratios dimensionless.
FVP-2  SCALE TYPE      Stevens' typology per term. Division and multiplication REQUIRE a
                    ratio scale with a true zero. If only ordinal data exist, say so:
                      the formula is a heuristic ordering, not an equation, and is
                            DEMOTED in print.
FVP-3  LIMITING CASES  Behaviour as each term -> 0 and -> infinity. Every divergence must
                    correspond to a real, nameable phenomenon or the formula is wrong.
FVP-4   FIXED POINTS      Solve for the invariant. State what it means.
FVP-5  IDENTIFIABILITY Can the parameters be recovered from obtainable data? If two distinc
                    parameter sets predict identical observations, the formula is not
                            testable as written.
FVP-6   FALSIFICATION     Preregister; fit on training; FREEZE; report held-
out; beat four
                                                  baselines  (persistence,  direct-
input, matched autoregressive,
                    SHUFFLED-INPUT control). State in advance what result kills it.
FVP-7  REPRODUCIBILITY Same operationalisation, different people, different laboratories:
                            do the parameters recover?
```

52.0.6 Why seven, and why in this order

Each gate kills a distinct class of error, and the classes do not overlap. A formula can be dimensionally perfect and scale-illegal. It can be scale-legal and unidentifiable. It can be identifiable and unfalsifiable. It can pass every one of the first six and still be a private fact about one person's rating habits, which is what FVP-7 is for.

Gate	Class of error it kills	Cost to run	What passing buys
FVP-1	Category error: adding unlike things	Hours	The equation is *about* something coherent
FVP-2	Operation error: arithmetic the data cannot support	Hours to years	The operations performed are meaningful
FVP-3	Structural error: behaviour with no referent	Hours	The formula's extremes describe real states
FVP-4	Vacuity: no invariant, therefore no content	Hours	The formula says something at equilibrium

Gate	Class of error it kills	Cost to run	What passing buys
FVP-5	Aliasing: distinct causes, identical predictions	Days	The parameters are recoverable in principle
FVP-6	Overfitting and post-hoc rescue	Months	The formula beat the alternatives on data it never saw
FVP-7	Idiosyncrasy: a result true of one person only	Years	The formula is about people, not about a person

The order is a cost ordering and it must be respected. Gates 1 through 4 are armchair work: they cost an afternoon and they kill most bad formulas. Gate 5 costs days. Gate 6 costs months. Gate 7 costs years and multiple laboratories. Running gate 6 on a formula that fails gate 1 is a way of spending a year to learn something an afternoon would have taught. In this manual's experience the ratio is stark: **the cheap gates kill most candidates, and the expensive gates are where the surviving candidates are actually decided.**

52.0.7 FVP-7 is the gate that decides, and the gate that gets skipped

Every one of the first six gates can be passed by a result true of exactly one person on exactly one instrument. A single subject, densely measured, can produce beautiful held-out prediction that beats all four baselines, and the whole thing can still be an artefact of how that one person uses the word "vivid."

FVP-7 is where the hard-science question is actually settled, and it is the gate the literature most often omits. It is omitted because it is expensive, because it requires coordination with people who do not share your priors, and because — the real reason — it is the gate most likely to return a negative. A framework that has passed gates 1 through 6 feels finished. It is not. The 36% figure is what the world looks like when a field routinely stops at gate 6.

Chapter 53 is FVP-2 in full, because scale type is where this manual's own central formula is most at risk. Chapter 54 is the deep form of FVP-7, because invariance across persons and instruments is what "reproducible measurement" actually means. Chapter 55 is FVP-6 in full. This chapter defines all seven and runs the complete protocol, end to end, on `M = (D/P)E`.

The Engineering

52.1 FVP-1 — Dimensions

Dimensional analysis is the cheapest true statement you can make about an equation, and it is astonishing how often it is skipped in the human sciences. It requires no data — only that you write down, beside each symbol, what a unit of that symbol *is*. Three rules, none negotiable:

```
R1  Only like-dimensioned quantities may be added or subtracted.
R2 The argument of ln, exp, sin, or any transcendental function must be dimensionless.
R3  Both sides of an equation must carry identical dimensions.
```

R2 does the most work in this corpus, because this corpus is logarithmic throughout. You cannot take the logarithm of three metres. You can take the logarithm of a *ratio* of metres, and the result is dimensionless — measured in nats, a unit of information, not of length.

Worked check: the surplus law.

```
u = 1 − (x . y)^2 = ||x ^ y||^2            x, y are UNIT vectors −
> u dimensionless, in [0,1]
N                                       a count −> dimensionless
1 + (N−1)u                            dimensionless + dimensionless −> OK by R1
F = ln(1 + (N−1)u)                  argument dimensionless −> OK by R2; F in nats
```

Passes. Note what the check enforces: x and y must be *unit* vectors, or u is not dimensionless and the construction fails at R2. That is the reason Paper A works on normalised states.

Worked check: the storage law.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t

Q      units of stored order     [q]
S      nats
eta    [q] per nat                        −> eta S has units [q]        OK
delta  dimensionless per period           −> (1−delta) Q has units [q]  OK
C      [q] per period                     −> subtractable from [q]      OK
h = eta ln N − delta Q − C              all three terms in [q]       OK by R1
```

Passes, and the check *earns something*: `eta` is not a pure number. **Eta is a conversion constant with units — the exchange rate between nats of interaction and units of durable structure.** Anyone reporting an eta without saying what `[q]` is has reported half a number. Chapter 4 said this in words; FVP-1 says it in units.

Worked check: the Moment equation.

```
M = (D / P) . E

For D/P to be dimensionless (R2 is not invoked here, but the interpretation demands it),
D and P must carry the SAME units.  Then:

r = D/P            dimensionless
M = r E            M inherits E's units exactly

Distortion: Delta = |ln r|    argument dimensionless −> OK by R2, Delta in nats
```

Passes, **conditionally** — and the condition is the whole content of the gate here. If perspective is measured in "confidence points" and perception in "cues detected per minute", the ratio is not dimensionless, it is a rate with units, and `M` does not inherit `E`'s units. The formula would then be false as written, not merely unmeasured.

FVP-1 therefore imposes a design requirement on the whole of Part VII: D and P must be instrumented on one common scale. That is a hard constraint and it is where most naive operationalisations of the Moment equation die. It is also the constraint that makes cross-modality matching (Chapter 53, §53.6) relevant rather than ornamental.

52.2 FVP-2 — Scale type

Chapter 53 is this gate, in full, and nothing here substitutes for it. The summary that belongs in the protocol definition:

```
NOMINAL    identity only        admissible transforms: any 1–1 relabelling
ORDINAL    order              admissible transforms: any strictly increasing g
INTERVAL   order + distance    admissible transforms: x -> a x + b,  a > 0
RATIO      order + distance
           + a true zero       admissible transforms: x -> a x,      a > 0
```

The rule that governs this manual: division and multiplication are meaningful only on a ratio scale. The reason is exactly the invariance structure above. A statement survives as a claim about the world only if it is unchanged under every transformation the scale admits. On a ratio scale, `g(D)/g(P) = (aD)/(aP) = D/P`: the ratio is invariant, so it carries information. On an interval scale, `(aD+b)/(aP+b)` depends on `b`, which was arbitrary. On an ordinal scale, `g(D)/g(P)` depends on the entire shape of `g`, and `g` was arbitrary.

FVP-2 returns one of three verdicts:

```
PASS     ratio-scaled with a demonstrated true zero -> the equation stands as an equation
DEMOTE    ordinal or interval only -> the formula is printed as a HEURISTIC ORDERING,
            with the demotion stated in the same paragraph as the formula
FAIL     the quantity has no defensible scale at all -> the term is not admissible
```

The DEMOTE verdict is the honest middle, and this manual uses it. A demoted formula is still useful: it can order occasions, flag directions, and structure attention. What it cannot do is carry a number that is then multiplied by something else. Chapter 53 §53.8 specifies exactly what survives demotion and what does not.

52.3 FVP-3 — Limiting cases

Take every term to zero and to infinity, one at a time, and read what the formula says. Then ask, for each result, **is there a real phenomenon with that name?**

This gate is a structure test, not a data test, and it is unreasonably effective. Formulas built by analogy almost always have a limit that is nonsense, and the nonsense stays invisible until someone takes the limit.

```
RULE FVP-3a  A divergence (a term going to infinity) must correspond to a real, nameable
         phenomenon. If it does not, the formula is wrong at the singularity, and
              almost always wrong well before it.

RULE FVP-3b  A zero must correspond to a real, nameable phenomenon. "It just goes to zero"
              is not an interpretation.

RULE FVP-3c  Where the limit is a genuine impossibility rather than a phenomenon, the
          formula must state its DOMAIN OF VALIDITY and stop there. Restricting the
             domain is legitimate. Pretending the limit is meaningful is not.
```

The Moment equation is the exemplar and it passes with unusual grace; §52.9 runs the numbers. The storage law also passes, and the passing is instructive: as `delta -> 0` the model gives `Q_t`

<= Q_0 + t eta ln N (Paper E, Theorem 3.3) — unbounded linear growth with no fixed point. That is not a defect; it is the exact description of a system with no forgetting, which is Chapter 3's subject, and which really does fail to converge. **The limit named a real phenomenon, so the gate passes.**

52.4 FVP-4 — Fixed points

Solve for the invariant. Two distinct objects go by this name and confusing them is a common error:

```
FIXED POINT OF A MAP      x* such that T(x*) = x*.  Dynamic. Banach guarantees existence and
                              convergence when T is a contraction.
                                                    Example:  Q_inf  =  (eta  F(u*)  -
c*) / delta   [Paper E, Thm 3.5]

INVARIANT OF AN EQUATION   the parameter configuration at which the equation reduces to an
                             identity. Static. No dynamics required.
                             Example: D = P  =>  M = E
```

Both must be reported. The map's fixed point tells you where iteration arrives; the equation's invariant tells you what the formula is *for*.

A formula with no invariant is usually vacuous. If no configuration of the parameters reduces it to something simpler and interpretable, it is a curve-fit rather than a law — laws are compressions, and compressions have degenerate cases.

FVP-4 also requires that the fixed point's *stability* be stated. `Q_inf` is attracting whenever `delta > 0`, with `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` — geometric convergence, Chapter 44. `D = P` is *not* dynamically attracting in the Moment equation as written, because the Moment equation is not a map: it is an instantaneous relation. Saying so is part of passing the gate. Nothing in `M = (D/P)E` makes `r` drift toward 1. That drift, if it exists, is a separate postulate about learning, and Chapter 35 states it separately and badges it.

52.5 FVP-5 — Identifiability

A formula is identifiable if distinct parameter values imply distinct observable predictions. If two different parameter sets predict exactly the same data, no amount of data will ever distinguish them, and the formula is not testable as written. This is a property of the *model*, not of the sample size, and it is checked algebraically before anything is measured.

Two kinds:

```
STRUCTURAL identifiability  distinct parameters -> distinct predictions, in principle,
                          with infinite noiseless data. An algebraic property.
PRACTICAL  identifiability  the parameters can be recovered with useful precision from the
                         data you can actually obtain. A statistical property.
```

Structural failure is fatal and no quantity of data repairs it. Practical failure can sometimes be fixed by better design.

The Moment equation has a structural identifiability problem, and it is severe. `M = (D/P)E` is invariant under two separate transformations, computed here on concrete values:

case	D	P	E	M = (D/P)E
A (reference)	8.0	4.0	3.0	**6.0000**
B (D,P both scaled by 2.5)	20.0	10.0	3.0	**6.0000**
C (r doubled, E halved)	16.0	4.0	1.5	**6.0000**
D (r quartered, E quadrupled)	2.0	4.0	12.0	**6.0000**

Four different worlds. One observation. **From M alone, none of D, P, E or even r is recoverable.** The first invariance (rows A and B) is benign in a sense — it says the ratio r is what matters and the common scale of D and P is free — but it means you cannot report D as a number without an anchor. The second invariance (rows A, C, D) is the dangerous one: **r and E trade off exactly, so a large distortion with a thin environment is observationally identical to undistorted contact with a rich one.**

This is not a fatal objection but a specification of what the instrument must supply. Three remedies, in ascending order of cost:

REMEDY I5-a ANCHOR the common scale. Fix a standard stimulus and define D and P relative to it. This removes the (cD, cP) freedom by convention. Cost: a standard.

REMEDY I5-b MEASURE E INDEPENDENTLY of M, by an instrument that does not pass through the observer. Physical measures of the environment; a second observer; a record. This breaks the r–E trade. Cost: an external instrument. THIS ONE IS ESSENTIAL.

REMEDY I5-c VARY the terms independently across occasions and use the TIME SERIES. If D and P and E move independently over time, the log-linear design matrix has full column rank and the exponents are recoverable. Cost: repeated measures, and the requirement that the terms genuinely decorrelate.

Remedy I5-c has a checkable condition, and it is the practical form of the gate. Writing the Moment equation in the log register (CANON_INTERNAL §1) gives

$$\ln M = b_0 + b_D \ln D + b_P \ln P + b_E \ln E + \text{noise}$$

with the theoretical prediction $(b_D, b_P, b_E) = (1, -1, 1)$. The parameters are recoverable if and only if the matrix of $[1, \ln D, \ln P, \ln E]$ has full column rank. **If D and P co-vary perfectly across occasions — if your frame always tracks your intake — the columns are collinear and no data set of any size will separate bD from bP.** That is a real risk, not a hypothetical one, because a well-adapted person's frame *does* track their intake. The design must therefore include occasions that break the coupling, and the analysis must report the collinearity diagnostics. A protocol that measures only ordinary days will fail FVP-5 in practice while passing it in principle.

52.6 FVP-6 — Falsification

The template is Paper E §5.3, and it is stated as a sequence because the order is what makes it work:

1. PREREGISTER Write the operationalisation, the model, the exponent targets, the

```
              baselines, the held-out split, and THE KILL CRITERIA. Deposit it
                   before collecting data. The kill criteria are the point.
2.  FIT             Estimate parameters on the training partition ONLY.
3.  FREEZE       Write the parameters down. They do not change again. Not for a better
                     fit, not for an outlier, not for a discovered confound.
4.  HELD-OUT      Apply the frozen parameters to data never used in fitting. Report the
                     out-of-sample statistic.
5.  BASELINES       Beat all four:
                       (i)   PERSISTENCE       yhat_t = y_{t-1}
                       (ii)  DIRECT-INPUT      the environment term alone
                  (iii) MATCHED AR      an autoregressive model of the same order
                  (iv)  SHUFFLED-INPUT   the same model, inputs permuted across time
6.  REPORT       Publish the failures alongside the successes, at the same prominence.
```

The shuffled-input control matters most and is most often absent. Every other baseline asks whether your model beats a simpler model. The shuffle asks a sharper question: *does the model's performance depend on the inputs being matched to the right occasions at all?* A model with enough free parameters, fitted to autocorrelated data, can post a respectable held-out score purely from the marginal distributions. Shuffling destroys the pairing while preserving every marginal. If the shuffled model does nearly as well as the real one, the real one learned nothing about the relation.

The kill criteria must be numeric and must be written before the data exist. For the Moment equation, this manual preregisters:

```
KILL-1  The 95% confidence interval for bP excludes -1.0 -> the DIVISION is not there.
KILL-2  The 95% confidence interval for bD excludes  1.0 -> the frame term is misspecified.
KILL-3  Held-out R^2 of the model fails to exceed the best baseline by at least 0.10.
KILL-4   Held-out R^2 of the SHUFFLED-INPUT control exceeds 0.20.
KILL-5     The  model  is  beaten  by  a  purely  additive  form  M  =  D  -
P + E on the same held-out set.
```

KILL-5 deserves comment because it is the specific alternative most likely to be true. Nothing about the phenomenology of perception forces a *multiplicative* combination. If frame and intake combine additively, the log-linear model will still fit reasonably — §52.9 computes exactly how reasonably — and only a direct comparison against the additive form exposes it. **Preregistering the rival functional form, not merely a null, is what separates a falsification test from a ritual.**

52.7 FVP-7 — Reproducibility

The gate: **same operationalisation, different people, different laboratories — do the parameters recover?**

There are three possible answers and they are not equally good, but only one of them is fatal:

```
OUTCOME A   UNIVERSAL CONSTANT
        The parameter takes the same value across persons, within sampling error.
         Strongest outcome. Rare. This is what a physical constant looks like.

OUTCOME B   PERSON-LEVEL CONSTANT
        The parameter varies between persons but is STABLE WITHIN a person across
```

```
        occasions and instruments. This is still measurement, and it is the normal
        case for psychological quantities. A person's height varies between persons
            too. What matters is that it is stable and that it predicts.

OUTCOME C   UNSTABLE
        The parameter varies within a person across occasions or instruments by as
        much as it varies between persons. FATAL. There is no quantity here; there
            is a fitting artefact.
```

Outcome B is a pass, and this must be said clearly, because the sceptic's demand for Outcome A is a demand physics itself would fail for most of its quantities. The mass of an object is not a universal constant. What makes mass a measurement is that the same object gives the same mass on different balances in different laboratories. The psychological analogue: the same person gives the same exponent on different instruments on different days. That is what Chapter 54 calls *specific objectivity*, and it is the real criterion.

The statistic that separates B from C is the between-person variance component. Given per-person estimates `b_i` with standard errors `se_i`:

```
tau^2  =  Var(b_i)  -  mean(se_i^2)          the true between-person variance
```

If `tau^2` is at or below zero, the parameter is a universal constant to within your precision (Outcome A). If `tau^2` is large and positive *and* the within-person estimates are stable across instruments, Outcome B. If `tau^2` is large and the within-person estimates are *not* stable across instruments, Outcome C, and the framework dies for that term.

The single most dangerous mistake at this gate is pooling. §52.9 computes the exact damage: pooling twelve heterogeneous subjects produced a confidence interval on `bD` about ten times narrower than the actual spread of the individuals' true values. The pooled analysis looked *more* precise the *more* wrong it was about every individual in it. Pooling is not a way of running FVP-7. It is a way of avoiding it.

52.8 The FVP row — the required print format

Every formula in Parts VI through XI carries this block. It is a specification artefact, not a summary: a reader must be able to check each cell against the chapter that produced it.

```
FVP ROW  --  <formula name>
  FVP-1  DIMENSIONS      <units of each term; verdict>
 FVP-2  SCALE TYPE    <scale required per term; scale available; PASS / DEMOTE / FAIL>
 FVP-3  LIMITING CASES  <each limit, its named phenomenon, domain of validity>
  FVP-4  FIXED POINTS    <invariant, meaning, stability>
  FVP-5  IDENTIFIABILITY <invariances; remedies required; rank condition>
 FVP-6  FALSIFICATION   <kill criteria; baselines; status: TESTED / UNTESTED>
  FVP-7  REPRODUCIBILITY <status: OUTCOME A / B / C / NOT YET RUN>
```

A gate that has not been run is reported as NOT YET RUN, never as blank and never as passed. This manual's own formulas mostly report NOT YET RUN at gates 6 and 7, and Chapter 56 is the costed programme for changing that. Printing the honest status is not an embarrassment; it is the only thing that makes the passed gates worth anything.

The Specification

Numbered, checkable requirements. A formulary entry satisfying V-1 through V-22 has been validated to this manual's standard. An entry that does not satisfy them is not a formula in this book; it is a note.

V-1 — No formula without an FVP row. The row appears in the same section as the formula, not in an appendix, not in a footnote. A formula printed without its row is a specification defect regardless of how well it performs.

V-2 — Units before symbols. Before any manipulation, tabulate every symbol with its dimension and its unit. Symbols whose unit cannot be stated are not admissible terms.

V-3 — Verify R1, R2 and R3 explicitly. State, in one line each: which additions are like-to-like, which transcendental arguments are dimensionless, and that both sides agree.

V-4 — Declare the common scale for every ratio. Any expression of the form X/Y must name the single scale on which both X and Y are measured, and the instrument that produces it.

V-5 — Classify each term on Stevens' typology, with evidence. The classification must cite the procedure that establishes it — magnitude estimation, cross-modality matching, ratio production, a Rasch calibration — not an assertion. "Treated as ratio" without a procedure is a FAIL, not a PASS.

V-6 — Demote in the same paragraph, not later. Where FVP-2 returns DEMOTE, the demotion appears in the paragraph that introduces the formula, in the same typeface, and the formula is thereafter written with an ordering relation rather than an equals sign wherever the text permits.

V-7 — Take every limit. For each term: the limit to zero, the limit to infinity, and the behaviour at any interior singularity. Report all of them, including the boring ones.

V-8 — Name the phenomenon at every limit. A limit reported without a named phenomenon fails FVP-3. If no phenomenon exists, restrict the domain and say where the formula stops applying.

V-9 — Report the invariant and its stability. State whether the fixed point is a fixed point of a map (and whether the map is a contraction) or an invariant of a relation (and therefore non-dynamic). Do not let the reader assume dynamics that the formula does not contain.

V-10 — Enumerate the invariance group. List every transformation of the parameters that leaves the predictions unchanged. This is FVP-5's core deliverable and it is algebra, not statistics. If the list is non-trivial, state the anchoring convention that removes each freedom.

V-11 — State the rank condition. For any formula to be estimated by regression, state the design condition under which the parameters are separable, and the collinearity diagnostic that will be reported.

V-12 — Preregister before data. Operationalisation, model, exponent targets, split, baselines, and kill criteria, deposited and timestamped. A preregistration written after the first look is not a preregistration; report it as an exploratory analysis and label it so.

V-13 — Freeze and never unfreeze. Parameters estimated on training data are written down before the held-out partition is touched. Any subsequent change to the model converts the study to exploratory and it must be relabelled throughout.

V-14 — All four baselines, always, including the shuffle. Persistence, direct-input, matched autoregressive, shuffled-input. Report all four scores whether or not the model beat them.

V-15 — Kill criteria are numeric and prior. No kill criterion may contain the words "materially", "substantially", or "clearly". A number and a comparison operator, or it is not a criterion.

V-16 — Preregister the rival functional form. At minimum, the additive analogue of any multiplicative formula and the multiplicative analogue of any additive one. Beating a null is not evidence for a specific functional form.

V-17 — Report per-subject parameters, never only the pool. Publish the individual estimates and their standard errors as a table. A pooled estimate may accompany them; it may not replace them.

V-18 — Compute and report `tau^2`. The between-person variance component, corrected for within-person sampling variance, with the FVP-7 outcome letter (A, B or C) stated.

V-19 — Test within-person stability across instruments, not only across occasions. Outcome B requires that the same person's parameter recovers on a *different* operationalisation. Occasion stability alone does not distinguish a measurement from a habit of responding.

V-20 — Report NOT YET RUN honestly. Any gate not executed is marked NOT YET RUN with a one-sentence statement of what running it would require. Blank cells are prohibited.

V-21 — Publish failures at equal prominence. A formula that failed a gate is reported in the chapter where it would have appeared, with the failure stated in the opening paragraph. Failed formulas are not silently removed between editions.

V-22 — Guard the protocol against itself. No FVP result may become a target (Goodhart; Campbell). Where an FVP statistic is used to select among candidate formulas, the selection must be disclosed and the winning formula must be re-run on fresh data before it is reported as validated.

The Failure Modes

Each entry: name, clinical description, damaged gate, diagnostic signature, detecting measurement, fix.

FM-52.1 . THE DRESSED-UP METAPHOR

Damages: FVP-1 and FVP-3.

An intuition is written with an equals sign and thereby acquires an unearned authority. The symbols have no units, the terms cannot be added, and no limit was ever taken. The formula is a diagram of a feeling.

Diagnostic signature: the author cannot say what a unit of any term is; every term is described qualitatively; taking a limit produces nonsense. A reliable tell: the formula never appears with numbers in it, anywhere, ever.

Detecting measurement: attempt V-2. Tabulate the units. The tabulation either completes or it does not.

Fix: either instrument the terms or rewrite the relation in words. **A well-written sentence is not a lesser thing than a badly-founded equation; it is a more honest one.**

FM-52.2 . THE UNANCHORED RATIO

Damages: FVP-5.

A ratio is reported as a number without any convention fixing the common scale, so the number means nothing across contexts. Two reports of `r = 1.4` from two instruments are not comparable and are nonetheless compared.

Diagnostic signature: `r` values drift systematically between studies, between instruments, or between the same person's sessions on different apparatus, with no substantive account of why.

Detecting measurement: run the same subject on two operationalisations in one sitting and regress one `r` on the other. A slope far from 1 with a high correlation is the signature: the quantity is real, the anchor is missing.

Fix: Remedy I5-a. Adopt and publish a standard stimulus. The choice of standard is arbitrary; the existence of one is not.

FM-52.3 . THE R-E TRADE

Damages: FVP-5, specifically.

The environment term is estimated from the same report that produced the moment, so `r` and `E` absorb each other and the model fits perfectly while identifying nothing. This is the single largest threat to the Moment equation and it is the reason Remedy I5-b is marked essential.

Diagnostic signature: excellent in-sample fit; `bE` and the intercept strongly correlated in the parameter covariance matrix; held-out performance collapses when the environment measure is replaced by an external one.

Detecting measurement: compute the correlation between estimated `ln r` and estimated `ln E` across occasions. Strong negative correlation is the trade appearing directly.

Fix: measure `E` by an instrument that does not pass through the observer. A physical record, a second observer, a logged event stream. Without this, the Moment equation is not testable.

FM-52.4 . THE MISSING SHUFFLE

Damages: FVP-6.

Three baselines are reported and the shuffled-input control is not. The model beats persistence, beats direct-input and beats an AR(1), and none of that establishes that it used the pairing between inputs and outcomes.

Diagnostic signature: the baselines reported are all *simpler models*, none is a *destroyed-pairing control*. The word "shuffle" or "permutation" does not appear in the methods.

Detecting measurement: run it. It costs one line of code and one refit.

Fix: run it, and report the number whatever it is. In the demonstration of §52.9, the real model scored `R^2 = 0.8091` held-out and the shuffled control scored `0.0059`. That gap is what a working model looks like. A gap of `0.81` versus `0.55` would have meant something very different.

FM-52.5 . THE POOLED FICTION

Damages: FVP-7, fatally, while appearing to strengthen it.

Individual subjects are pooled into one regression. The pooled confidence interval is narrow because `n` is large. The narrowness is reported as precision. It is not precision; it is a statement about the *average* of a population, presented as though it were a statement about the *members*.

Diagnostic signature: per-subject estimates are not reported; `n` in the results is the number of observations rather than the number of subjects; the confidence interval is implausibly tight for a psychological quantity.

Detecting measurement: fit every subject separately and compute `tau^2` (V-18). In §52.9's heterogeneous case, `tau = 0.2534` against a pooled standard error of `0.0257`. **The pooled interval was about ten times narrower than the real spread.**

Fix: V-17 and V-18. Publish the per-subject table. Report `tau`. If the target is a group-level claim, use a model with an explicit random effect and report the variance component, not only the mean.

FM-52.6 . THE SILENT DEMOTION

Damages: FVP-2, and everything downstream of it.

Ordinal data are collected, arithmetic that requires a ratio scale is performed on them, and the scale question is never raised. This is the most common failure in the entire human-sciences literature and it is invisible because it is universal.

Diagnostic signature: seven-point ratings entering a multiplication or a division; means and standard deviations of Likert items; ratios of Likert means reported to two decimal places.

Detecting measurement: Chapter 53's re-labelling test. Apply an arbitrary strictly increasing transformation to the raw responses, recompute the result, and see whether the conclusion survives. If it does not, the conclusion was about the labels.

Fix: build a real scale (Chapter 53, §53.5-53.7) or demote in print (V-6). There is no third option, and the pretence that there is one is what the 36% figure is partly made of.

FM-52.7 . THE EPICYCLE

Damages: FVP-6.

A prediction fails and the model acquires a term. The new term is justified post hoc and is never itself preregistered. Repeat, and the model becomes unfalsifiable by accretion rather than by design.

Diagnostic signature: the model has more free parameters than it had last year; every added parameter was added after a specific failure; there is no held-out test that postdates the last addition.

Detecting measurement: count parameters against the preregistration. Compare held-out performance of the current model to the *original* frozen model on *new* data.

Fix: freeze, and re-run the full FVP-6 suite on fresh data with the amended model preregistered. An amended model is a new model and inherits none of the old one's credit.

FM-52.8 . GOODHART CAPTURE OF THE PROTOCOL

Damages: all seven, by making them a target.

The FVP becomes a compliance exercise. Formulas are constructed to pass the gates rather than tested by them; kill criteria are written loose enough to survive; the shuffle control is run on a version of the data with the pairing partially preserved.

Diagnostic signature: every formula in the corpus passes every gate. This is the tell. **A protocol that nothing fails is measuring nothing** (Goodhart; Campbell).

Detecting measurement: the pass rate. In an honest corpus it is well below one.

Fix: V-21 and V-22. Publish failures. Re-run selected formulas on fresh data. And keep the pass rate itself on the dashboard, because a rising pass rate is the earliest signal of capture.

Worked Example

The complete protocol, run end to end on the Moment equation, with every number computed.

```
M = (D / P) . E        D perspective, P perception, E environment, M moment
```

The empirical portions (FVP-6 and FVP-7) are run here on **simulated data**, and that label is load-bearing. **No claim is made here about human beings.** What is demonstrated is the *protocol*: what the analysis looks like when the formula is true, what it looks like when it is false, and what the numbers have to be for a reader to be entitled to believe it. Chapter 56 costs out the programme that would run this on people.

52.9.1 FVP-1 — Dimensions

term	meaning	dimension	unit
`D`	perspective / frame strength	`[s]`	one common subjective magnitude scale, anchored
`P`	perception / intake strength	`[s]`	the same scale, same anchor
`r = D/P`	calibration ratio	dimensionless	pure number
`E`	environment magnitude	`[e]`	environmental units, externally instrumented
`M`	moment magnitude	`[e]`	inherits `E` exactly
`Delta = |ln r|`	distortion	dimensionless	nats

R1: no unlike additions occur. R2: `ln r` takes a dimensionless argument. R3: left side `[e]`, right side `dimensionless x [e] = [e]`.

VERDICT: PASS, conditional on `D` and `P` sharing one anchored scale. The condition is not cosmetic; it forces the instrument design and it is the reason Chapter 53 §53.6 matters.

52.9.2 FVP-2 — Scale type

`D` and `P` are divided, so both require a ratio scale with a true zero. `E` is multiplied, so it too requires a ratio scale. `M` inherits.

The true zeros are interpretable, which is a necessary condition and a good sign: `P = 0` is total absence of intake, `D = 0` is total absence of frame, `E = 0` is total absence of environment. Each is a state one can describe. **A quantity whose zero cannot be described does not have a true zero, and the ratio scale is unavailable.**

Whether ratio-scaled `D` and `P` can actually be *built* is Chapter 53's subject. The honest verdict for this manual, today:

VERDICT: DEMOTE, pending instrumentation. Magnitude estimation gives a validated route (Chapter 53, §53.5), and the route has produced ratio scales for sensory magnitudes for seventy years. It has not yet been applied to `D` and `P` by anyone. Until it has, `M = (D/P)E` is printed in this manual as a **heuristic ordering with a demonstrated structure**, not as a measured equation. Everything below shows what passing would look like; it does not show that we have passed.

52.9.3 FVP-3 — Limiting cases

With `D = 6`, `E = 2`, sweeping `P` down:

P	r = D/P	M = rE	Delta = |ln r|
4.0000	1.50	3.00	0.4055
2.0000	3.00	6.00	1.0986
1.0000	6.00	12.00	1.7918
0.5000	12.00	24.00	2.4849
0.2500	24.00	48.00	3.1781
0.1000	60.00	120.00	4.0943
0.0500	120.00	240.00	4.7875
0.0100	600.00	1200.00	6.3969
0.0010	6000.00	12000.00	8.6995

With `P = 3`, `E = 2`, sweeping `D` down:

D	r	M	Delta
12.0000	4.0000	8.0000	1.3863
6.0000	2.0000	4.0000	0.6931
3.0000	**1.0000**	**2.0000**	**0.0000**
1.5000	0.5000	1.0000	0.6931
0.7500	0.2500	0.5000	1.3863

D	r	M	Delta
0.3000	0.1000	0.2000	2.3026
0.0300	0.0100	0.0200	4.6052
0.0030	0.0010	0.0020	6.9078

The four limits and their named phenomena (CANON_INTERNAL §4):

```
P -> 0   M -> infinity   frame with no intake    -> hallucination, dogma, apophenia
D  ->  0       M  ->  0                   intake  with  no  frame          -
> no moment; sensation is not experience
E -> 0    M -> 0           no environment           -> no moment
D = P     M  = E           THE FIXED POINT          -> undistorted contact
```

Every limit names a real, describable state, and the divergence names the most clinically familiar one of the four. A frame that receives no correcting intake and grows without bound is the structure of dogma, of the closed ideology, and of the psychotic elaboration. The formula's singularity is not a mathematical embarrassment to be excused; it is the formula's most recognisable prediction.

Note also the shape of the `Delta` column. Distortion grows *logarithmically* while `M` grows *linearly* in `r`. Cutting intake by a factor of ten multiplies the moment by ten but adds only `ln 10 = 2.3026` nats of distortion. **In the log register, catastrophic over-framing is a modest number, and that is exactly the register the person inside it experiences.** The lens theorem (Chapter 36) applied to its own pathology.

Domain of validity: the formula is stated for `D >= 0`, `P > 0`, `E >= 0`. At `P = 0` the expression is undefined and the model is replaced by the statement "unbounded", not by a number.

VERDICT: PASS.

52.9.4 FVP-4 — Fixed points

Set `M = E` and solve: `(D/P)E = E` for `E > 0` gives `D/P = 1`, hence `D = P`.

The invariant is `D = P`, and its meaning is veridical contact: the moment equals the environment. Frame matches intake exactly and adds nothing and subtracts nothing. This is the formula's most important content and it is the reason the formula is worth stating at all — it provides an equation for a condition usually described only in adjectives.

Stability: **the Moment equation is a relation, not a map, and `D = P` is therefore not dynamically attracting.** Nothing in `M = (D/P)E` drives `r` toward 1. Any claim that calibration improves with experience is a separate postulate about learning, is not contained in this formula, and is badged separately in Chapter 35. Reporting this is part of passing the gate: **a static invariant presented as though it were an attractor is a real and common misreading, and it is the author's job to forestall it.**

Contrast with the storage law, where the fixed point *is* dynamic: `Q_inf = (eta F(u*) - c*)/delta` (Paper E, Thm 3.5), approached geometrically with `|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|`, guaranteed by Banach because `delta > 0` makes the map a contraction.

VERDICT: PASS, with the non-dynamic status stated.

52.9.5 FVP-5 — Identifiability

The invariance group, from §52.5: `(D, P, E) -> (cD, cP, E)` for any `c > 0`, and `(D, P, E) -> (kD, P, E/k)` for any `k > 0`. The four-row table in §52.5 exhibits both.

VERDICT: FAIL as written; PASS under Remedies I5-a, I5-b and I5-c jointly. All three are required, and the specification for any study of the Moment equation must include all three. This is a genuine defect of the formula in its bare form and printing it as a failure is the point of having the gate.

The rank condition (V-11): the design matrix `[1, ln D, ln P, ln E]` must have full column rank across occasions. In the demonstration series below, the terms were generated with independent innovations and moderate autocorrelation, so the condition holds. **In a real study it may not, and the collinearity between `ln D` and `ln P` must be reported before any exponent is interpreted.**

52.9.6 FVP-6 — Falsification, run

Preregistered: a single-subject series of `T = 200` occasions; `D`, `P`, `E` each log-normally distributed with first-order autocorrelation (0.55, 0.55, 0.40 respectively, so that persistence and AR baselines are genuine competitors, not straw men); log-normal multiplicative noise on `M` with `sigma = 0.35` in the log register. Model:

```
ln M = b0 + bD ln D + bP ln P + bE ln E + e
targets: (b0, bD, bP, bE) = (0, 1, −1, 1)
```

Split: occasions 1-140 train, 141-200 held out. Parameters frozen after training.

Fitted on training (n = 140), then frozen:

parameter	estimate	s.e.	95% CI	target	t vs target
b0	-0.0295	0.1191	[-0.2629, 0.2039]	0.0	-0.248
bD	1.0205	0.0504	[0.9217, 1.1193]	1.0	0.407
bP	-0.9505	0.0627	[-1.0734, -0.8275]	-1.0	0.790
bE	0.8985	0.0725	[0.7563, 1.0407]	1.0	-1.399

All four intervals contain their targets. KILL-1 and KILL-2 are not triggered.

Held out (n = 60), frozen parameters, against all four baselines:

model	held-out R^2	RMSE
the Moment model	**0.8091**	**0.4067**
persistence (`yhat_t = y_{t−1}`)	-0.3322	1.0744
direct-input (E alone)	0.3797	0.7331
matched AR(1)	0.0756	0.8950
shuffled-input control	**0.0059**	**0.9281**

The margin over the best baseline is `0.8091 − 0.3797 = 0.4294`, comfortably past KILL-3's threshold of 0.10. The shuffled control returns `0.0059`, far below KILL-4's ceiling of 0.20, and its

fitted coefficients collapse toward zero (`bD = 0.1249`, `bP = −0.0524`, `bE = 0.2024`) — which is what destroying the pairing is supposed to do.

The persistence baseline's *negative* `R^2` is not an error: it means the baseline predicts worse than the training mean, the correct behaviour for persistence against a series whose innovations dominate its autocorrelation.

Now the same suite run against two false worlds — the part that makes the suite worth running.

Counterfactual A — the division is not there. Data generated from `M = D.E`, with `P` present in the record but causally inert. Same fitting procedure:

```
bD = 1.0890   CI [0.9448, 1.2332]
bP = −0.0709  CI [−0.2009, 0.0590]     <−− KILL−1 TRIGGERED
bE = 0.9017   CI [0.7485, 1.0549]
```

The `bP` interval misses `−1.0` by `0.799`. **The protocol kills the formula, correctly, and it kills it on the specific term that was wrong** — which is more than a global goodness-of-fit statistic would have told anyone.

Counterfactual B — combination is additive, not multiplicative. Data generated from `M = D − P + E` in raw units. Fitting the log-linear multiplicative model:

```
log−linear fit to additive data:   bD = 1.6477   bP = −1.9381   bE = 1.2322
held−out R^2 of the WRONG multiplicative model  = 0.6758
held−out R^2 of the correct ADDITIVE model      = 0.9694
additive coefficients: 0.930, −0.882, 0.977  (targets 1, −1, 1)
```

This is the important result in the chapter. The wrong functional form still achieved a held-out `R^2` of `0.68`. Against persistence, against direct-input, against AR(1) it would have won handsomely. A study that preregistered only a null would have reported the multiplicative Moment equation as confirmed. Only KILL-5 — the preregistered *rival functional form* — exposes it, and it exposes it decisively: `0.9694` against `0.6758`.

The exponents also announce the problem to anyone who looks: `bD = 1.65` and `bP = −1.94` are nowhere near 1 and −1, so KILL-1 and KILL-2 both fire as well. **A formula that predicts its exponents is far more falsifiable than one that merely predicts a direction, and that is an argument for stating exponent targets in the preregistration.**

Power. From the observed `s.e.(bP) = 0.0627` at `n = 140`, the standard error scales as `0.7419 / sqrt(n)`. Required sessions to detect a given deviation from the exponent target:

deviation from target	n for 80% power	n for 90% power
0.50	18	24
0.40	27	37
0.30	48	65
0.25	70	93
0.20	108	145
0.15	192	258
0.10	432	579
0.05	1728	2314

Read this table before designing the study, not after. Detecting a `0.2` deviation in an exponent needs about 110 sessions from one person; detecting `0.1` needs about 430. A thirty-session pilot can only rule out gross misspecification, and should say so.

VERDICT on the demonstration: PASS on data generated from the model; correctly FAILS on both counterfactuals. VERDICT for the Moment equation as applied to human beings: NOT YET RUN.

52.9.7 FVP-7 — Reproducibility, run

Twelve simulated subjects, 200 occasions each, same operationalisation.

Case 1 — homogeneous population. All twelve subjects generated with true exponents `(1, −1, 1)`.

subject	bD (s.e.)	bP (s.e.)	bE (s.e.)	covers (1, -1, 1)
1	0.950 (0.051)	-0.969 (0.051)	0.997 (0.052)	YYY
2	1.053 (0.055)	-1.024 (0.056)	1.088 (0.050)	YYY
3	1.011 (0.061)	-0.946 (0.046)	0.979 (0.059)	YYY
4	0.964 (0.050)	-1.111 (0.055)	1.071 (0.058)	YnY
5	1.019 (0.053)	-0.937 (0.055)	0.976 (0.060)	YYY
6	1.111 (0.059)	-1.007 (0.052)	1.075 (0.057)	YYY
7	0.951 (0.045)	-0.955 (0.047)	1.048 (0.053)	YYY
8	1.033 (0.046)	-1.026 (0.052)	1.079 (0.059)	YYY
9	1.034 (0.052)	-1.048 (0.050)	1.100 (0.055)	YYY
10	0.977 (0.044)	-1.038 (0.048)	0.985 (0.053)	YYY
11	1.013 (0.049)	-1.028 (0.046)	0.986 (0.066)	YYY
12	0.992 (0.058)	-1.036 (0.058)	0.966 (0.058)	YYY

```
mean bD = 1.0092  sd 0.0464       coverage of bD on 1.0  : 12/12
mean bP = −1.0104 sd 0.0504       coverage of bP on −1.0 : 11/12
mean bE = 1.0291  sd 0.0517       coverage of bE on 1.0  : 12/12

bD: between−subject variance 0.00216 ; mean within−subject variance 0.00273
    tau^2 = max(0, 0.00216 − 0.00273) = 0.00000   −>  tau = 0.0000
```

`tau` is zero to numerical precision: the between-subject spread is entirely accounted for by sampling error. **This is OUTCOME A — a universal constant, to within the precision of the study.** The single miss at subject 4 is the expected behaviour of a 95% interval across twelve tests; one miss in twelve is not evidence of heterogeneity, and reporting it as such would be an error in the other direction.

Case 2 — heterogeneous population. True exponents drawn per subject with standard deviation 0.30 around `(1, −1, 1)`.

subject	true (bD, bP, bE)	estimated bD (s.e.)	estimated bP (s.e.)	estimated bE (s.e.)	covers
1	(0.759, -1.397, 0.925)	0.723 (0.050)	-1.302 (0.053)	0.895 (0.050)	nnn
2	(1.126, -0.659, 1.033)	1.102 (0.052)	-0.697 (0.052)	1.104 (0.053)	YnY
3	(0.834, -1.235, 1.225)	0.842 (0.054)	-1.261 (0.049)	1.327 (0.060)	nnn
4	(1.490, -0.918, 0.630)	1.490 (0.060)	-0.920 (0.053)	0.633 (0.048)	nYn
5	(0.713, -0.520, 1.061)	0.716 (0.044)	-0.452 (0.049)	1.094 (0.054)	nnY
6	(0.480, -1.025, 0.651)	0.563 (0.048)	-1.058 (0.054)	0.768 (0.057)	nYn
7	(0.811, -1.146, 0.786)	0.825 (0.047)	-1.117 (0.049)	0.804 (0.055)	nnn
8	(1.166, -1.019, 0.823)	1.120 (0.058)	-0.932 (0.054)	0.863 (0.053)	nYn
9	(1.123, -0.751, 0.507)	1.124 (0.055)	-0.757 (0.046)	0.465 (0.051)	nnn
10	(0.923, -1.294, 0.948)	0.876 (0.045)	-1.332 (0.054)	1.052 (0.052)	nnY
11	(0.613, -0.994, 0.989)	0.653 (0.056)	-1.055 (0.050)	1.016 (0.059)	nYY
12	(0.909, -1.314, 0.881)	0.919 (0.053)	-1.305 (0.051)	0.908 (0.049)	YnY

```
mean bD = 0.9128  sd 0.2587       coverage of bD on 1.0  : 2/12
mean bP = -1.0157 sd 0.2767       coverage of bP on -1.0 : 4/12
mean bE = 0.9106  sd 0.2306       coverage of bE on 1.0  : 5/12

bD: between-subject variance 0.06693 ; mean within-subject variance 0.00270
    ratio 24.78   ->   tau^2 = 0.06423   ->   tau = 0.2534
```

Every subject's own exponents are recovered accurately — compare the true and estimated

columns. **The estimates are excellent and the universal constant is absent.** Heterogeneity is roughly 25 times the sampling noise. This is **OUTCOME B**: the exponent is a person-level parameter, not a species-level one.

Outcome B is a pass provided V-19 is also satisfied — provided the same person's exponent recovers on a *different instrument*. If it does, the exponent is that person's and the framework has a measurement. If it does not, the exponent belongs to the instrument and the framework has nothing.

Now the damage from pooling. The same heterogeneous data, fitted as one pooled regression over all 2400 observations:

```
bD = 0.8474   s.e. 0.0257   CI [0.7969, 0.8979]
bP = -1.0451  s.e. 0.0257   CI [-1.0955, -0.9947]
bE = 0.9094   s.e. 0.0267   CI [0.8571, 0.9617]
```

The pooled bP interval contains −1.0, so a pooled analysis would have reported the division confirmed. Meanwhile:

```
tau (true spread of individual bD)        = 0.2534
pooled standard error on bD               = 0.0257
ratio                                     = 9.9
```

The pooled confidence interval is about ten times narrower than the actual spread of the individuals it describes. It is not a wrong interval — it is a correct interval about the wrong quantity. It describes the population mean; it was read as describing a person. That substitution is FM-52.5, it is extremely common, and it is one of the mechanisms behind the replication gap this chapter opened with. **An estimate can be simultaneously precise, unbiased about its own target, and useless for every individual in the sample.**

52.9.8 The completed FVP row

```
FVP ROW  --  THE MOMENT EQUATION   M = (D/P).E

 FVP-1  DIMENSIONS    D,P share one anchored scale [s]; r dimensionless; M inherits [e];
                         Delta = |ln r| in nats. R1, R2, R3 all verified.
                         PASS, conditional on the common scale.

 FVP-2  SCALE TYPE    D, P, E all require RATIO scale with true zero. All three zeros are
                         describable states. Instrumentation not yet built.
                   DEMOTE -- printed in this manual as a heuristic ordering with
                    demonstrated structure, pending Chapter 53's construction.

 FVP-3  LIMITING CASES  P->0 unbounded (dogma/apophenia); D->0 zero (sensation without
                  experience); E->0 zero (no environment); D=P gives M=E (contact).
                         All four named. Domain: D>=0, P>0, E>=0.
                         PASS.

 FVP-4  FIXED POINTS   Invariant D=P, meaning veridical contact. STATIC, not an attractor;
                        the formula contains no dynamics driving r toward 1.
                         PASS, with non-dynamic status stated.
```

```
  FVP-5  IDENTIFIABILITY Two invariances: (D,P,E)->(cD,cP,E) and (D,P,E)-
>(kD,P,E/k).
                       Bare formula is NOT identifiable. Requires Remedy I5-
a (anchored
                           standard), I5-b (independent measurement of E --
ESSENTIAL), and
                    I5-c (independent variation across occasions, with the rank
                          condition reported).
                          FAIL as written; PASS under all three remedies.

 FVP-6  FALSIFICATION  Kill criteria KILL-1..KILL-5 preregistered, including the rival
                  additive form. Protocol demonstrated on simulated data: passes on
                  true data (held-out R^2 = 0.8091 vs shuffle 0.0059), correctly
                          kills both counterfactuals. Power table computed.
                          NOT YET RUN on human data.

 FVP-7  REPRODUCIBILITY Protocol demonstrated; Outcome A and Outcome B both exhibited;
                          pooling damage quantified at a factor of ~10.
                          NOT YET RUN on human data.
```

Two of the seven gates report NOT YET RUN and one reports DEMOTE. That is the honest state of this manual's central formula, and it is printed here, in the chapter that defines the standard, so that no reader can mistake the ambition for the achievement.

The ambition is nonetheless achievable, and that is the chapter's closing claim. The demonstration above is arithmetic on a design that could be run starting Monday with a hundred and forty sessions and a preregistration. What stands between this manual and a validated Moment equation is Chapter 53's scale construction, Chapter 55's falsification suite, and Chapter 56's costed programme — not a philosophical obstacle, and not the nature of subjectivity.

Psychology and sociology are not pseudoscience. They are sciences with a measurement problem and a replication problem. Formulas do not fix that. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does — and that is what the next two chapters are for.

Chapter 53 . Why Division Requires a Ratio Scale

The Law

This is the crux of Volume Two. Every other chapter in Parts VII, VIII and IX either survives this one or is demoted by it.

The Moment equation divides. `M = (D/P)E` takes perspective, divides it by perception, and multiplies the result by environment. **Division and multiplication are not universally available operations. They are licensed by the scale type of the quantities involved, and they are licensed only on a ratio scale with a true zero.** If `D` and `P` are not ratio-scaled, the expression `D/P` does not fail to be accurate — it fails to be *about anything*, because its value can be set to whatever you like by a relabelling that changes nothing in the world.

That is the claim of this chapter, and it is not a matter of taste. It is a theorem about invariance, and the theorem is proved in three lines.

53.0.1 Stevens' typology

S. S. Stevens, *On the theory of scales of measurement* (*Science*, 1946), classified measurement scales by **the group of transformations under which the scale's information is preserved.** A scale is not a set of numbers; it is a set of numbers *together with the class of relabellings that leave what the numbers mean unchanged.* That class is the scale.

```
SCALE      WHAT IS PRESERVED           ADMISSIBLE TRANSFORMATIONS g
--------   ------------------------   ------------------------------------
---------
NOMINAL     identity / distinctness      any one-to-one map            x -
> g(x), g injective
ORDINAL   identity + order         any strictly increasing g   x -> g(x), g' > 0
INTERVAL    + equality of differences    positive affine                x -
> a x + b,  a > 0
RATIO     + a true zero            similarity              x -> a x,     a > 0
ABSOLUTE   + the unit itself           identity only               x -> x
```

Each row's transformation group is a *subgroup* of the row above it. Nominal admits the most relabellings and therefore carries the least information; absolute admits none and carries the most. Moving down the table is buying information by giving up freedom.

53.0.2 The meaningfulness criterion

A statement made in terms of scale values is meaningful if and only if its truth value is unchanged under every admissible transformation of the scale.

This is the operative definition and it is the whole of the chapter's logic. It is not a philosophical preference. It follows from what a scale *is*: if `g` is admissible, then `g(x)` and `x` encode the same empirical facts by definition, so any statement true of one and false of the other is a statement about the labelling, not about the world.

The criterion is mechanical to apply. Write the statement. Apply an admissible `g`. See if it survives.

53.0.3 The division theorem

Theorem 53.1. *Let X and Y be measured on a common scale with admissible transformation group G. The statement "X/Y = k" is meaningful if and only if G is the similarity group {g(x) = ax : a > 0} or a subgroup of it — that is, if and only if the scale is ratio or absolute.*

Proof. Three cases.

```
RATIO.     g(x) = a x,  a > 0.
           g(X)/g(Y) = (aX)/(aY) = X/Y.
           The ratio is invariant under every admissible g. MEANINGFUL.

INTERVAL.  g(x) = a x + b,  a > 0.
           g(X)/g(Y) = (aX + b)/(aY + b).
           Take X = 20, Y = 10. Under g = identity the ratio is 2.000.
           Under g(x) = 1.8x + 32 the ratio is 68/50 = 1.360.
         The value depends on b, which the scale leaves free. NOT MEANINGFUL.

ORDINAL.   g strictly increasing, otherwise arbitrary.
       g(X)/g(Y) can be given ANY value greater than 1 (for X > Y > 0) by a suitable
       choice of g. Demonstrated constructively in section 53.4. NOT MEANINGFUL.
```

The converse direction is immediate: if `g(X)/g(Y) = X/Y` for all admissible g and all positive `X`, `Y`, then taking `Y` fixed shows g is linear through the origin, so `G` is contained in the similarity group. QED.

Therefore: `D/P` is meaningful if and only if `D` and `P` are ratio-scaled with a true zero. There is no third possibility, no partial credit, and no way to get there by having a lot of data. A million observations on an ordinal scale is a million ordinal observations.

53.0.4 What an interval scale does license

The theorem is often misread as saying interval scales are useless. They are not. **On an interval scale, ratios of *differences* are meaningful, and this is exactly the claim the scale supports.** Verified on the values `x = (10, 20, 35, 40)` under four admissible affine transformations:

transformation	values	`(x2−x1)/(x4−x3)`	`x2/x1`
identity	10, 20, 35, 40	**2.000000**	2.000000
`g(x) = 1.8x + 32`	50, 68, 95, 104	**2.000000**	1.360000
`g(x) = 2x − 7`	13, 33, 63, 73	**2.000000**	2.538462
`g(x) = 0.4x + 100`	104, 108, 114, 116	**2.000000**	1.038462

The ratio of differences is rock stable at 2.000000 under every admissible relabelling. The ratio of values takes four different answers. That contrast is the whole distinction, and it is why interval-scaled temperature supports "the rise from Monday to Tuesday was twice the rise from Wednesday to Thursday" but not "Tuesday was twice as warm as Monday."

53.0.5 The badge

Note on the status of Stevens' typology. Stevens' four-way classification is a **proposal about scientific practice**, not a theorem, and it has been criticised — most sharply by Joel Michell, whose *Measurement in Psychology: A Critical History of a Methodological Concept* (1999) argues that psychology largely assumed its attributes were quantitative rather than testing whether they are. The representational theory of measurement (Krantz, Luce, Suppes & Tversky, *Foundations of Measurement*, 1971) puts the whole subject on axiomatic footing and treats Stevens' types as special cases.

None of that weakens Theorem 53.1, which is an invariance argument and holds within any of these frameworks. Whatever you call the scale, if the admissible relabellings include anything beyond `x -> ax`, ratios of values are not invariant, and a formula that divides is reporting a property of the labelling. The badge is on the *typology*. The theorem stands unbadged.

The Engineering

53.1 The four scales, in operational detail

NOMINAL. Numbers as names. Jersey numbers, diagnostic codes, room identifiers, the numbering of Book One's twelve nodes.

```
Permissible:    counting members of a class; mode; chi-square; contingency tables;
                  entropy of the partition; "same or different"
NOT permissible:  order, mean, median, difference, ratio, correlation
```

A nominal scale still carries real information — the entropy of a partition is a genuine quantity, and Chapter 7's lineage census (EC-9) uses exactly that. **You can do serious quantitative work on nominal data; you cannot do arithmetic on the labels.**

ORDINAL. Order without spacing. Rankings, hardness scales, seniority, most rating scales as they are actually used, "how anxious do you feel, 1 to 7."

```
Permissible:    median; percentiles; rank correlation (Spearman, Kendall); order statistic
                  monotone regression; "greater than"; sign tests
NOT permissible:  mean; standard deviation; difference; ratio; Pearson correlation;
                  any arithmetic on the labels
```

INTERVAL. Equal differences, arbitrary zero. Celsius, Fahrenheit, calendar dates, standardised scores, and — importantly for Chapter 54 — properly constructed Rasch measures in logits.

```
Permissible:    mean; standard deviation; Pearson correlation; differences; ratios of
                  DIFFERENCES; linear regression; t-tests and ANOVA
NOT permissible:  ratios of VALUES; "twice as much"; multiplication of two measures;
                  proportional statements
```

RATIO. Equal differences plus a non-arbitrary zero denoting genuine absence. Length, mass, duration, count, Kelvin, frequency in hertz, and — the existence proof this volume depends on — Stevens' magnitude estimates of sensory intensity.

```
Permissible:    everything above, plus ratios of values, multiplication, division,
```

```
geometric mean, coefficient of variation, logarithms of the values,
      "twice as loud"
```

Note the last two entries. **The geometric mean and the logarithm are ratio-scale operations, and this entire manual runs on both.** `ln N_eff`, `ln M`, `|ln r|`, the octave, the nat: every one of them presupposes ratio-scaled inputs. That is not an incidental fact about the notation. **The logarithmic register of Volume Two (CANON_INTERNAL §1) is available only because ratio scales support it, and this chapter is the price of admission.**

53.2 What a true zero actually is

"True zero" is not "the smallest number on the form". It is a specific empirical claim:

`A scale has a TRUE ZERO if the value 0 denotes the complete ABSENCE of the attribute, and that absence is an empirically identifiable state rather than a chosen origin.`

Three tests, in ascending order of strength:

Test Z1 — describability. Can you describe the state the zero denotes, and would a second person recognise it? Zero mass is describable. Zero degrees Celsius is not a state of absence; it is the freezing point of water, chosen for convenience. **This test is necessary and cheap and almost nobody runs it.**

For the Moment equation's terms, the zeros pass Z1 and they pass it well:

```
P = 0   total absence of intake. Sensory deprivation, dissociation, the closed eye.
D = 0   total absence of frame. Raw sensation with no interpretive structure at all.
E = 0    total absence of environment. No situation to have a moment about.
```

Test Z2 — additivity / concatenation. Is there an operation that concatenates two instances of the attribute such that the scale values add? Two one-kilogram masses on a pan read two kilograms. This is the classical criterion and it is why length, mass and duration are unproblematic. For subjective quantities there is generally no physical concatenation operation, which is precisely Michell's objection and the historical reason psychology's measurement claims were contested.

Test Z3 — ratio consistency. Do judgements of ratio behave consistently? If an observer says `B` is twice `A` and `C` is twice `B`, do they say `C` is four times `A`? Does halving compose with halving to give quartering? **This is the test that replaces concatenation for subjective quantities, and it is the test magnitude estimation is built on.** It is weaker than Z2 but it is empirical, it is checkable, and it either holds or it does not.

A scale that passes Z1 and Z3 is admissible as ratio for this manual's purposes, with the provenance stated. That is the standard adopted here, it is the standard psychophysics adopted, and it is badged: Z3 is a behavioural consistency criterion, not a concatenation, and a critic who holds out for Z2 is not being unreasonable.

53.3 What goes wrong on an interval scale

Temperature is the canonical demonstration because everyone already knows the answer and almost nobody carries the lesson into their own data.

Celsius pair	C ratio	Fahrenheit values	F ratio	Kelvin values	K ratio
20 / 10	2.000	68.0 / 50.0	1.360	293.15 / 283.15	1.0353
25 / 5	5.000	77.0 / 41.0	1.878	298.15 / 278.15	1.0719
15 / -5	-3.000	59.0 / 23.0	2.565	288.15 / 268.15	1.0746
10 / 0	**undefined**	50.0 / 32.0	1.562	283.15 / 273.15	1.0366

Read the rows across. **The same physical pair of temperatures produces the ratio 2.000, 1.360 and 1.0353 depending on nothing but the labelling convention.** Row three produces a *negative* ratio in Celsius and a positive one in Fahrenheit — the two labellings disagree about the sign. Row four is undefined in Celsius and perfectly finite in Fahrenheit.

Only the Kelvin column is a real quantity, because Kelvin has a true zero and its admissible transformations are similarities alone. **The ratio 1.0353 is a fact about the world. The ratio 2.000 is a fact about Anders Celsius.**

Now transpose this. A questionnaire scored 1 to 7 has its zero at "1", chosen because forms usually start at one. A scale scored 0 to 6 measuring the same thing has its zero at "0". These two instruments differ by `g(x) = x − 1`, a perfectly admissible affine transformation. A `D/P` of `6/3 = 2.0` on the first instrument is `5/2 = 2.5` on the second. **Nothing changed but the printer's choice, and the reported distortion moved by 25%.**

53.4 What goes wrong on an ordinal scale — the demonstrations

Ordinal is worse than interval, and it is worse in a way that surprises people who have absorbed the temperature lesson. On an interval scale, at least the *family* of possible ratios is constrained by a two-parameter group. **On an ordinal scale the admissible transformations form an infinite-dimensional family, and the ratio can be sent anywhere.**

Demonstration 1 — the magnitude is arbitrary. Four occasions of a frame/intake report on a seven-point instrument, `D = (6, 3, 5, 2)`, `P = (3, 6, 4, 4)`. Apply admissible strictly increasing relabellings and recompute `D/P`:

relabelling g	D/P across the four occasions	ranking, most over-framed first
`g(x) = x` (raw)	2.000, 0.500, 1.250, 0.500	1, 3, then 2 and 4 TIED
`g(x) = x^2`	4.000, 0.250, 1.562, 0.250	1, 3, then 2 and 4 TIED
`g(x) = ln(x+1)`	1.404, 0.712, 1.113, 0.683	1, 3, **2, 4**
`g(x) = x + 10`	1.231, 0.812, 1.071, 0.857	1, 3, **4, 2**
`g(x) = e^(x/2)`	4.482, 0.223, 1.649, 0.368	1, 3, **4, 2**
`g(x) = 2^x`	8.000, 0.125, 2.000, 0.250	1, 3, **4, 2**
bespoke monotone	20.000, 0.050, 12.500, 0.500	1, 3, **4, 2**

The last column is the alarming one. Under three admissible relabellings occasions 2 and 4 are tied, under one occasion 2 is worse, and under four occasion 4 is worse. **The tie is broken in both directions by relabellings that changed nothing about what the respondent reported.**

Occasion 1's distortion is reported as anything from 1.231 to 20.000 — a factor of sixteen — with no change to a single response. **Every one of these relabellings is admissible. The instrument recorded the same thing in every row.**

Demonstration 2 — the ordering across occasions reverses. Take `D = (7, 4)` and `P = (4, 2)`:

relabelling g	r occasion 1	r occasion 2	verdict
`g(x) = x`	1.7500	2.0000	occasion 2 more over-framed
`g(x) = x^2`	3.0625	4.0000	occasion 2 more over-framed
`g(x) = x^3`	5.3594	8.0000	occasion 2 more over-framed
`g(x) = ln(x+1)`	1.2920	1.4650	occasion 2 more over-framed
`g(x) = 2^x`	8.0000	4.0000	**occasion 1 more over-framed**
`g(x) = e^x`	20.0855	7.3891	**occasion 1 more over-framed**

Which occasion was more distorted is not a fact about the person. It is a fact about the shape of the response labels. This is the finding that matters clinically, because "which session was worse" is exactly the question a practitioner asks.

Demonstration 3 — the constructive result. For `D > P > 0` on an ordinal scale, *any* target ratio can be achieved by an admissible relabelling. With `D = 7`, `P = 4`, use `g(x) = x^k`, which is strictly increasing on the positives for every `k > 0`, giving `g(D)/g(P) = (7/4)^k`:

target ratio	required k	achieved
1.05	0.0872	1.0500
1.50	0.7245	1.5000
2.00	1.2386	2.0000
5.00	2.8760	5.0000
10.00	4.1146	10.0000
100.00	8.2292	100.0000

A quantity you can set to any value you like by relabelling is not a measurement. It carries exactly zero bits of information beyond the ordering it started with.

Demonstration 4 — the ratio can become negative or undefined. `g(x) = x − 2.5` is strictly increasing and therefore admissible on an ordinal scale. Applied to the four occasions of Demonstration 1, occasion 4 has `D = 2 -> −0.5` and `P = 4 -> 1.5`, so `r = −0.333` and `Delta = |ln r|` is undefined over the reals. **An admissible relabelling turned the distortion measure into a non-number.** No amount of careful statistics downstream repairs this.

53.5 Building a ratio scale: magnitude estimation

The good news, and it is very good news, is that ratio scales for subjective quantity are not hypothetical. Stevens' method of magnitude estimation produces them, it has produced them

since 1957, and it replicates (CANON_INTERNAL §3).

The procedure, specified.

```
ME-1  Present stimuli from the continuum in IRREGULAR order, spanning at least 1.5 orders
       of magnitude. A narrow range biases the exponent toward 1.
ME-2  Instruct: "Assign any positive number to the first stimulus. Assign numbers to the
     rest so that the numbers stand in the same RATIO as the sensations. If one seems
       three times as strong, give it three times the number."
ME-3  Impose NO upper bound and NO fixed modulus. The observer's own first number is the
       modulus and it is free.
ME-4  Normalise each observer by their own geometric mean across stimuli. This removes the
       arbitrary modulus without touching any ratio.
ME-5  Aggregate across observers by the GEOMETRIC mean, never the arithmetic mean.
ME-6   Regress log(response) on log(intensity). The slope is the exponent a.
ME-7   Report the exponent with its interval, the log-log R^2, and the per-
observer spread.
```

The prediction being tested is Stevens' power law:

```
psi = k I^a       ->       log psi = log k + a log I
```

A straight line in log-log is not a curve-fitting convenience. **It is the falsifiable content of the claim: if the observer's judgements were not ratio-consistent, the log-log plot would bend, and it does not.**

Worked recovery. Ten simulated observers, seven intensities spanning 1 to 64, generated with a true exponent of 0.67, independent free moduli (the largest observer's modulus was 31.15 times the smallest), and 18% log-normal judgement noise.

```
geometric mean across observers, by intensity:
   0.2300   0.4124   0.5753   0.9469   1.7176   2.6619   4.2328

log-log fit on the geometric means:
   a_hat = 0.6987   s.e. 0.0165   95% CI [0.6664, 0.7310]      true a = 0.67
   log-log R^2 = 0.99723
```

The exponent is recovered, and the log-log linearity is essentially perfect. Per-observer exponents, each fitted separately:

```
0.737  0.569  0.694  0.653  0.755  0.774  0.720  0.631  0.728  0.726
mean 0.6987   sd 0.0631
```

Note what this is: FVP-7 Outcome B in miniature. The exponent is a person-level parameter with a real spread of about 0.06, not a single universal number, and that is the normal and acceptable state of a psychological measurement.

Why the geometric mean is mandatory (ME-5). With free moduli, the arithmetic mean is dominated by whichever observer happened to pick large numbers. In the run above, one observer alone contributed **27.5%** of the arithmetic mean at the lowest intensity and **21.1%** at the highest, in a panel of ten where each should carry 10%. Under the geometric mean each observer contributes exactly 1/n of the log, by construction. **The arithmetic mean of magnitude estimates**

silently weights observers by their choice of modulus, which is the one thing the method declares arbitrary.

53.6 The ceiling: why a rating scale cannot do this

Here is the computation that should end the practice of treating slider responses as ratio data. The same magnitude-estimation simulation — same true exponent `0.67`, same observers, same noise — run four times, changing only the response format:

response format	recovered exponent	bias
unbounded, free modulus (proper ME)	0.6678	-0.0022
0-100 slider	0.5763	-0.0937
0-50 slider	0.4076	-0.2624
1-7 rating	**0.0114**	**-0.6586**

A seven-point rating scale recovered an exponent of 0.011 from a process whose true exponent is 0.67. The measured quantity was destroyed almost entirely — not degraded, destroyed. The bound compresses the top of the range, the compression is worst where the signal is strongest, and the log-log slope collapses.

This is not an argument that rating scales are useless. They are excellent ordinal instruments. **It is an argument that a bounded response format cannot produce a ratio scale, ever, and that arithmetic performed on its output is arithmetic on the bound.** A ratio scale must be unbounded above, because doubling must always be possible; a scale with a ceiling has a largest number, and "twice that" has no representation.

53.7 Cross-modality matching — the validation that makes it science

Magnitude estimation on its own could, in principle, be measuring how people use numbers rather than how they experience intensity. **Cross-modality matching is the test that rules this out, and it is the strongest piece of evidence in the whole psychophysical corpus.**

The procedure uses no numbers at all. The observer adjusts one continuum — squeezes a dynamometer, turns a light up — until it *feels* as intense as a stimulus presented on a different continuum. The exponent of that matching function is then a prediction with **no free parameters**:

```
If  psi_1 = k_1 I_1^a   and   psi_2 = k_2 I_2^b,
then matching psi_1 = psi_2 gives  I_2 proportional to I_1^(a/b).
```

```
The MATCHING EXPONENT is a/b -- fully determined by two exponents measured in two
separate experiments that never involved matching.
```

The full prediction matrix from this manual's canonical exponents (CANON_INTERNAL §3). Rows are the continuum the observer *adjusts*; columns are the continuum *presented*; entries are the predicted matching exponent `a/b`:

adjusted pre-sented	Loudness 0.67	Bright 0.33	Length 1.00	Salt 1.30	Heavy 1.45	Warmth 1.60	Shock 3.50
Loudness (0.67)	1.0000	0.4925	1.4925	1.9403	2.1642	2.3881	5.2239
Brightness (0.33)	2.0303	1.0000	3.0303	3.9394	4.3939	4.8485	10.6061
Visual length (1.00)	0.6700	0.3300	1.0000	1.3000	1.4500	1.6000	3.5000
Taste, salt (1.30)	0.5154	0.2538	0.7692	1.0000	1.1154	1.2308	2.6923
Heaviness (1.45)	0.4621	0.2276	0.6897	0.8966	1.0000	1.1034	2.4138
Warmth (1.60)	0.4188	0.2062	0.6250	0.8125	0.9062	1.0000	2.1875
Electric shock (3.50)	0.1914	0.0943	0.2857	0.3714	0.4143	0.4571	1.0000

Every cell is a testable prediction. And the matrix carries an internal consistency requirement that is itself a falsification test — **transitivity**:

```
match loudness -> heaviness   = 0.4621
match heaviness -> brightness = 4.3939
product                       = 2.0303
direct loudness -> brightness = 2.0303       AGREES

match brightness -> warmth    = 0.2062
match warmth -> loudness      = 2.3881
product                       = 0.4925
direct brightness -> loudness = 0.4925       AGREES
```

This is what a ratio scale looks like from outside. Three experiments, no shared parameters, and the third one's answer is forced by the first two. When that holds, the scale is not a story about how people use numbers; it is a measurement, because a measurement is exactly a set of numbers whose relations compose.

Honest note, required by CANON_INTERNAL §3: the canonical exponents are nominal, they vary with method, range and instruction, and they are properly reported as *constants of a procedure* rather than constants of nature. Empirically, cross-modality matching exponents recover the predicted a/b well but not perfectly, and the residual discrepancies are a known and studied feature of the literature. **The transitivity test is a real test that can fail, which is what makes passing it worth something.**

53.8 Ratio production — the third instrument

Magnitude estimation asks for a number given a stimulus. **Ratio production inverts it: give the observer a number and a standard, and have them produce the stimulus.** "Adjust this until it is half as bright." "Set this until it is three times as heavy."

Production is the strongest of the three because it is the direct behavioural test of Z3, and it composes:

```
RP–1   halve, then halve again -> should equal a single quartering
RP–2   double, then double again -> should equal a single quadrupling
RP–3   triple, then take a third -> should return to the standard
RP–4   the production function and the estimation function should be inverse,
       giving exponents that are reciprocals
```

RP-4 is the sharpest check and it is a genuine falsification opportunity: estimation and production are known to yield systematically slightly different exponents (the "regression effect", in which observers under-use whichever variable they control). **Report both, report the discrepancy, and take the geometric mean of the two exponents if a single figure is needed.** Concealing the discrepancy would be the error; it is a known property of the method and it is documented.

53.9 The other route: additive conjoint measurement

There is a second path to metric structure and it deserves a full hearing, because it is the path that rescues ordinal data — and it is the path Chapter 54 walks.

Luce & Tukey, *Simultaneous conjoint measurement* (*Journal of Mathematical Psychology*, 1964), proved a remarkable thing. **Suppose you have only an ORDINAL ordering of outcomes, but the outcomes are produced by two factors crossed with each other. If that ordinal table satisfies certain purely ordinal axioms, then interval scales for both factors exist and are essentially unique.** The metric is not assumed. It is *derived from order*, and its existence is a testable consequence of the ordering.

The central axiom is **double cancellation**:

```
if   (a1, b2) >= (a2, b1)   and   (a2, b3) >= (a3, b2)
then (a1, b3) >= (a3, b1)
```

Every symbol is an ordinal comparison. Nothing is added or multiplied. Verified on two 3x3 tables, all 36 antecedent-satisfying triples tested in each:

```
Table A = [[1, 2, 4], [2, 4, 8], [4, 8, 16]]    (a_i * b_j: additive in logs)
   36 triples tested, 0 violations       -> additive structure available

Table B = [[1, 2, 4], [2, 9, 3], [4, 3, 16]]     (non-additive)
   36 triples tested, 3 violations       -> additive structure REFUTED
```

This is the honest answer to "but all I have is ordinal data". You may still have metric structure — but you have to earn it by testing the cancellation conditions, not by assuming them. Double cancellation, together with solvability and an Archimedean condition, is what licenses the interval scale. **Michell's objection is answered not by assertion but by running the test.**

The deterministic form above is brittle: one noisy observation produces a violation. The probabilistic form of the same idea — which tolerates noise, and which is the reason Chapter 54 exists — is the Rasch model. **Rasch measurement is additive conjoint measurement with error, and the person-by-item table is the two-factor crossing.** That is the connection between these two chapters and it is not decorative: it is the same theorem, made usable.

53.10 The honest demotion — what to do when you only have ordinal data

Most of the time, in most real settings, you will have ordinal data. This section specifies exactly what you may then do, and it is the section a practitioner should photocopy.

What SURVIVES demotion. Given D and P recorded on the *same* ordinal scale with the *same* labels:

```
SURVIVES-1  The SIGN of the comparison. If D > P on the raw labels, then g(D) > g(P) for
         every strictly increasing g. So "this occasion was over-framed" is an
           admissible ordinal claim, PROVIDED both terms use one instrument.
SURVIVES-2  Within-occasion direction of change over time, if the instrument is fixed.
          "The gap narrowed between March and June" is ordinal and admissible.
SURVIVES-3  Rank correlations. Spearman and Kendall are invariant under monotone
             relabelling by construction.
SURVIVES-4  Medians, percentiles, order statistics, sign tests, monotone regression.
SURVIVES-5  The QUALITATIVE structure of the formula: the limiting cases, the fixed point
         D = P as a describable condition, and the three-regime classification
           (over-framed / calibrated / under-framed). These are order facts.
```

What DOES NOT survive.

```
DIES-1   The magnitude of r. Demonstration 3: settable to anything.
DIES-2   The ordering of r ACROSS occasions. Demonstration 2: reverses.
DIES-3   Delta = |ln r| as a number. Demonstration 4: can become undefined.
DIES-4  Any arithmetic that feeds r into another formula. In particular, M as an input
             to the self-storage law Q_{t+1} = (1-delta)Q_t + eta M_t -
C_t is NOT available
       from ordinal D and P. The join between the two central formulas requires
           ratio scales.
DIES-5   Means and standard deviations of the raw responses.
DIES-6   Pearson correlations of the raw responses.
```

DIES-4 is the expensive one. **The join between the Moment equation and the storage law (CANON_INTERNAL §8) is the architectural keystone of Volume Two, and it is available only on ratio scales.** That is stated here, in the chapter that establishes it, and not buried.

The Spearman/Pearson demonstration, on forty paired ordinal responses, computed under four admissible relabellings:

relabelling	Pearson r	Spearman rho
identity	0.6474	**0.6692**
g(z) = z^2	0.5167	**0.6692**
g(z) = e^(z/2)	0.4370	**0.6692**

relabelling	Pearson r	Spearman rho
g(z) = ln(z+1)	0.7223	**0.6692**

Pearson moves from 0.44 to 0.72 — a range wide enough to change any conclusion — while Spearman is identical to four decimal places in every row, as it must be. **If your data are ordinal, the rank statistic is not a conservative choice. It is the only correct one.**

The print convention (FVP-2 DEMOTE, per V-6). A demoted formula is printed with an ordering relation and an explicit statement, in the same paragraph:

```
DEMOTED FORM:

    Over-framed  <=>  D > P          Calibrated  <=>  D = P
    Under-framed <=>  D < P

  On ordinal D and P the classification above is admissible. The magnitude D/P is NOT,
  and no value of D/P from this instrument may be multiplied, averaged, or entered into
    any downstream formula.
```

This is not a defeat. It is a considerable amount of usable structure — three named regimes, a fixed point, a direction of correction — obtained honestly and stated in a way a critic can check. What it is not is an equation, and it is not printed as one.

The Specification

S-1 — Classify every term before any arithmetic. For each symbol, record: the scale type, the instrument that produced it, and the transformation group that instrument admits. A term whose group cannot be stated is not admissible.

S-2 — Apply the meaningfulness test explicitly. For every statement the analysis will make, apply an admissible g and check invariance. Record the check. This costs minutes.

S-3 — Never divide or multiply below ratio. No exception, no "approximately ratio", no "close enough for practical purposes". Theorem 53.1 has no epsilon in it.

S-4 — Establish the true zero by Z1 and Z3, and say which. Describe the state the zero denotes in one sentence. Report the Z3 ratio-consistency check with its numbers. If Z2 concatenation is available, use it and say so; if not, say that too.

S-5 — Both terms of a ratio on one instrument. D and P must come from the same scale with the same anchor. Two instruments, however similar, do not license division of one by the other without a demonstrated common calibration.

S-6 — Response formats for ratio data must be unbounded above. No sliders, no fixed maxima, no "1 to 10". The ceiling table in §53.6 is the reason: a seven-point format recovered 0.011 where the truth was 0.67.

S-7 — Free modulus, always. The observer chooses their own first number. Anchoring the first response imposes a scale and destroys the ratio property being measured.

S-8 — Span at least 1.5 orders of magnitude in the stimulus. Narrow ranges bias exponents toward 1 and can manufacture apparent linearity out of nothing.

S-9 — Irregular presentation order. Monotone sequences invite response strategies and produce sequential dependencies that masquerade as the function being measured.

S-10 — Aggregate by geometric mean. Never the arithmetic mean of magnitude estimates. §53.5 quantifies the damage: one observer of ten carrying 27.5% of the weight.

S-11 — Report the log-log `R^2`. Linearity in log-log is the falsifiable content of the power law. A curved log-log plot is a finding and must be reported as one, not smoothed.

S-12 — Report per-observer exponents and their spread. The panel mean is not the result. The distribution is the result, and it determines the FVP-7 outcome letter.

S-13 — Validate by cross-modality matching. Where two continua have independently measured exponents, run the matching experiment and compare against the predicted `a/b`. This is the test that distinguishes measurement from number-usage habit.

S-14 — Run at least one transitivity triangle. Three continua, three matching experiments, and the composition check of §53.7. Report the agreement, or the failure.

S-15 — Run ratio production as a cross-check. RP-1 through RP-4. Report the estimation and production exponents separately and report their discrepancy; do not silently average.

S-16 — Where only ordinal data exist, test double cancellation before assuming metric structure. Report the number of antecedent-satisfying triples tested and the number of violations. A table with violations does not support an additive structure.

S-17 — Demote in print, in the same paragraph. V-6 of Chapter 52. The demotion appears with the formula, in the same typeface, not in a footnote and not in an appendix.

S-18 — Use rank statistics on ordinal data. Spearman or Kendall, never Pearson; median, never mean. §53.10 quantifies the alternative.

S-19 — Never feed a demoted quantity into another formula. DIES-4. A heuristic ordering cannot be an input to the storage law, cannot be averaged across people, and cannot be differenced across time as a number.

S-20 — State the scale provenance beside every reported number. The form is: `r = 2.14 [ratio scale, magnitude estimation, free modulus, geometric-mean aggregation, n = 14 observers, a_hat = 0.71 (0.04)]`. A number without provenance is not admissible in this manual.

The Failure Modes

FM-53.1 . THE LIKERT METRE

Damages: FVP-2, and everything downstream.

Seven-point responses are treated as though the intervals between them were equal and the bottom of the scale were a true zero. Means are computed, ratios are taken, and the results are reported to two decimal places.

Diagnostic signature: a mean of 4.73 on a 1-7 scale, reported without a distributional summary; a ratio of two such means; a standard deviation of ordinal responses; any Pearson correlation on raw Likert items.

Detecting measurement: the relabelling test. Recompute the headline result under `g(x) = x^2` and `g(x) = 2^x`. If the conclusion moves, it was a conclusion about the labels.

Fix: build a real scale (§53.5-53.8), earn an interval scale by conjoint or Rasch analysis (§53.9, Chapter 54), or demote (§53.10). **These are the only three options and the third one is respectable.**

FM-53.2 . THE ARBITRARY ZERO

Damages: FVP-2 specifically at the ratio boundary.

An interval scale is treated as ratio because its numbers happen to be positive. Positivity is not a true zero. Fahrenheit is positive across most of the inhabited world and is not a ratio scale.

Diagnostic signature: the author cannot describe the state that "0" denotes; the zero coincides suspiciously with the bottom of a form, the start of a study, or a sample mean.

Detecting measurement: Test Z1. Ask what the zero *is*. Then shift the origin by one unit and see whether any reported ratio survives.

Fix: locate a describable zero and re-anchor, or restrict all claims to differences and ratios of differences (§53.0.4), which interval scales genuinely support.

FM-53.3 . THE MODULUS LEAK

Damages: the aggregation step of a magnitude-estimation study.

Magnitude estimates are averaged arithmetically. The observers who chose big numbers dominate. The group exponent becomes a weighted statistic where the weights are the very quantity the method declared arbitrary.

Diagnostic signature: group means much larger than the median observer's values; the group exponent shifting when a single observer is removed; a modulus range across observers exceeding a factor of five with no normalisation step reported.

Detecting measurement: compute each observer's share of the group mean at each stimulus level. Under correct aggregation each carries `1/n`. In §53.5's run, arithmetic aggregation gave one of ten observers 27.5%.

Fix: S-7 and S-10. Normalise by the observer's own geometric mean, then aggregate geometrically.

FM-53.4 . THE CEILING

Damages: FVP-2, invisibly and catastrophically.

The response format has an upper bound. The upper part of the stimulus range compresses against it. The recovered exponent is biased toward zero, and the bias grows as the bound tightens.

Diagnostic signature: responses piling up at the maximum; a log-log plot that is straight at low intensities and bends flat at high ones; a recovered exponent implausibly close to zero.

Detecting measurement: the proportion of responses at the maximum. Any nonzero proportion is disqualifying for ratio claims. Compare the exponent fitted on the lower half of the range to the exponent fitted on the whole.

Fix: S-6. Remove the bound. If the bound cannot be removed — some field instruments cannot — the data are ordinal and must be demoted. **The 1-7 row of the table in §53.6 is what happens otherwise: the true exponent 0.67 recovered as 0.011.**

FM-53.5 . THE SILENT PROMOTION

Damages: all of FVP-2, and the reader's ability to check.

Ordinal data are promoted to ratio by fiat somewhere in the middle of an analysis pipeline, usually at the point where a mean is taken, and the promotion is never stated.

Diagnostic signature: the methods section describes a rating scale; the results section reports ratios; no sentence connects the two. The word "assume" does not appear.

Detecting measurement: trace every reported number back to its instrument. The promotion happens at a specific line and it can be found.

Fix: S-1 and S-20. Provenance beside every number. If the promotion is deliberate, badge it as a postulate with its justification and its consequences, per CANON §2.

FM-53.6 . THE CROSS-INSTRUMENT DIVISION

Damages: FVP-1 and FVP-2 jointly.

`D` is measured by one instrument and `P` by another, and the two are divided. Even if both are ratio scales, their units differ, so the quotient is a rate with units rather than a dimensionless calibration ratio, and the fixed point `D = P` has no meaning.

Diagnostic signature: the two terms come from different sections of a protocol, different response formats, or different modalities; `r` clusters implausibly far from 1 across all subjects, which is the units mismatch showing.

Detecting measurement: the distribution of `r` across a healthy sample. A systematic offset in `ln r` common to everyone is a unit conversion, not a population-wide distortion.

Fix: S-5. One instrument, one anchor, or an explicit cross-modality calibration (§53.7) converting one to the other with the conversion exponent reported.

FM-53.7 . THE UNTESTED CANCELLATION

Damages: FVP-2 via the conjoint route.

Additive conjoint measurement or a Rasch analysis is invoked to justify treating ordinal data as metric, and the cancellation conditions are never actually tested. The theorem is cited; its hypothesis is not checked.

Diagnostic signature: "we used a Rasch model, so the measures are interval" with no fit statistics, no invariance test, no violation count.

Detecting measurement: S-16 for the deterministic case; Chapter 54's fit statistics and split-sample invariance test for the probabilistic one.

Fix: run the test. **A conjoint or Rasch analysis that does not report its fit is a promotion by fiat wearing a better suit.**

Worked Example

Constructing a ratio scale for P, perceptual intake, and then delivering the FVP-2 verdict on D/P. This is a design, costed in effort, that a competent laboratory could run; the numbers below that come from simulation are labelled as such.

53.11.1 The attribute and its zero

P is defined as **the magnitude of registered sensory and situational intake during a bounded episode.** The zero state, per Test Z1: complete absence of intake — an episode in which nothing whatever was registered. This is describable, it is recognisable, and observers agree on instances of near-approach to it (dreamless sleep, deep absorption in an unrelated task, dissociative episodes). **Z1 passes.**

Z2 concatenation is unavailable: there is no operation that physically joins two intakes so their magnitudes add. **Z2 not available, and this is stated, not glossed.** The scale will rest on Z3.

53.11.2 The instrument

```
Stimulus continuum:  recorded episodes of controlled richness, spanning a factor of at
                least 40 in a pre-measured environmental index (count of distinct
                 events, distinct speakers, distinct objects manipulated -- an
                     EXTERNAL measure, satisfying Remedy I5-b of Chapter 52).
Task:             immediately after each episode, magnitude estimation of "how much
                  you took in", free modulus, unbounded, per ME-1 through ME-7.
Order:                 irregular, per S-9.
Aggregation:           geometric, per S-10.
```

Note that the external environmental index does double duty. It is the stimulus variable for the power-law fit here, and it is the independent measure of E that Chapter 52's FVP-5 declared essential. **One instrument, two gates.**

53.11.3 Expected analysis, demonstrated on simulated data

Using the §53.5 simulation as the template — ten observers, seven levels, free moduli spanning a factor of 31, 18% log-normal noise:

```
a_hat (geometric-mean aggregation)  = 0.6987  s.e. 0.0165  95% CI [0.6664, 0.7310]
log-log R^2                          = 0.99723
per-observer exponents               = 0.737 0.569 0.694 0.653 0.755
```

```
                                        0.774 0.720 0.631 0.728 0.726
per-observer mean 0.6987, sd 0.0631
```

What would have to be true for this to count. Not that the exponent takes a particular value — nothing in this manual predicts the exponent of perceptual intake. What must be true is:

```
C1  the log-log relation is straight (R^2 above about 0.98, no systematic curvature)
C2   no responses at any ceiling, because there is no ceiling
C3  ratio production composes: halve-then-halve equals quarter, within tolerance (Z3)
C4  cross-modality matching against an established continuum recovers the predicted a/b
C5   the same observer's exponent recovers on a second occasion, weeks apart
```

C1 and C2 are properties of a single session. C3 is Z3, the ratio-scale warrant. **C4 is the one that makes it science**, because it is a prediction from two other experiments with no free parameters. C5 is FVP-7 within a person, and it is the precondition for Outcome B.

53.11.4 The same construction for `D`

`D`, perspective, is defined as **the magnitude of interpretive frame brought to the episode** — how much structure the observer supplied. Its zero: raw sensation with no interpretation, which is describable and, at the limit, is what the Moment equation's `D -> 0` case names.

Critically, `D` must be estimated on the same instrument, in the same session, with the same modulus, per S-5. The observer gives one number for intake and one number for frame, both referred to the same freely chosen first number. This is what makes them commensurable and it is what makes `D/P` dimensionless (Chapter 52, FVP-1).

This is a real design constraint with a real cost: the observer must hold one modulus across two judgements, which is harder than holding it across one continuum, and the resulting scale will be noisier. **The alternative — two instruments, two anchors, one division — is FM-53.6 and is not admissible.**

53.11.5 The verdict

Assembling everything:

```
FVP-2 -- SCALE TYPE -- THE MOMENT EQUATION  M = (D/P).E

  Term D:  RATIO required (it is divided).
           Zero describable: YES (Z1 passes).
           Concatenation available: NO (Z2 unavailable; stated).
           Ratio-consistency established: NOT YET (Z3 not run).
           Instrument built: NO.
        STATUS: ratio scale is CONSTRUCTIBLE by the design in 53.11.2-53.11.4;
                    NOT YET CONSTRUCTED.

  Term P:  identical status to D.

  Term E:  RATIO required (it multiplies).
           Zero describable: YES. External physical index available: YES.
           STATUS: ratio scale AVAILABLE today via the environmental index.
```

```
  VERDICT: DEMOTE.

   Until D and P are instrumented per 53.11.2 and validated per C1-
C5, the Moment equation
  is printed in this manual in its DEMOTED FORM (section 53.10):

     Over-framed <=> D > P ;  Calibrated <=> D = P ;  Under-framed <=> D < P

  The three-regime classification, the four limiting cases, and the fixed point D = P are
  ADMISSIBLE on ordinal data from a single common instrument. The magnitude D/P, the
  distortion Delta = |ln r|, and the join to the storage law via M_t are NOT, and are
   reported nowhere in Volume Two as measured quantities.
```

53.11.6 What this chapter concludes

The crux was never whether subjective quantities can be measured on a ratio scale. Psychophysics settled that affirmatively seventy years ago, and the transitivity table in §53.7 is what the settlement looks like. The crux is whether *these particular* subjective quantities — frame and intake — have been so measured. They have not. The route is specified, the instrument is designed, the validation criteria C1 through C5 are written down, and the work has not been done.

That is a smaller and far more useful claim than the one a less careful book would make, and it has the property that it can be discharged. **Ratio-scale measurement of `D` and `P` is a finite programme of laboratory work, not a philosophical impasse**, and Chapter 56 costs it out.

Until then, the demotion stands, in print, in the same paragraph as the formula, every time. That is the discipline the Formula Validation Protocol exists to impose, and this chapter is the reason it has teeth.

Chapter 54 . Specific Objectivity

The Law

Chapter 53 established what a ratio scale is and why D/P needs one. It ended with a construction — magnitude estimation, cross-modality matching, ratio production — and with the honest verdict that the construction has not been carried out for perspective and perception.

This chapter is the second route, and it is the one that answers the harder question. A ratio scale built by magnitude estimation is still a scale built by *one instrument on one occasion*. Chapter 52's FVP-7 asks something more: **does the measure survive a change of instrument and a change of population?** That is the real criterion of hard measurement, and there is a body of theory that addresses it directly, achieves it, and can be checked.

54.0.1 The requirement, stated by Rasch

Georg Rasch, *Probabilistic Models for Some Intelligence and Attainment Tests* (1960), and more explicitly in *On general laws and the meaning of measurement in psychology* (1961), set out the condition he called **specific objectivity**:

The comparison of two persons must be independent of WHICH ITEMS were used to compare them.
The comparison of two items must be independent of WHICH PERSONS were used to compare them.

Read those two lines slowly, because they are more demanding than they look. The first says that if instrument A and instrument B both measure the same attribute, then A's verdict on whether Alice exceeds Bob must agree with B's — not approximately, not on average, but as a structural property of the model, holding for every pair of persons and every pair of items.

This is the property physical measurement has and psychological measurement usually does not. A kitchen scale and a laboratory balance agree about which of two objects is heavier, and they agree about *by how much* once the units are matched. They agree regardless of what else was weighed that day. Two questionnaires purporting to measure the same construct routinely disagree, and the disagreement is normally reported as "modest convergent validity" rather than as the failure of measurement it is.

Specific objectivity is precisely the property Volume Two needs, because the internal/external correspondence (CANON_INTERNAL §5) claims that the same parameters describe a mind and a city. That claim is empty unless the parameters mean the same thing when measured by different instruments at different scopes. **Invariance is not a nice-to-have property of a good measure. Invariance is what "measure" means.**

54.0.2 The model

For a dichotomous response X_ni of person n to item i:

```
                     exp(theta_n − beta_i)
P(X_ni = 1) = ──────────────────────────────
                1 + exp(theta_n − beta_i)

theta_n   the person's location on the latent continuum   (logits)
beta_i    the item's location on the same continuum       (logits)
```

Equivalently, and more revealingly:

```
        P(X_ni = 1)
ln  -------------  =  theta_n - beta_i
        P(X_ni = 0)
```

The log-odds of success is the DIFFERENCE between a person parameter and an item parameter, and nothing else. No item discrimination, no guessing parameter, no interaction term. That austerity is the entire point: it is what buys separability, and every parameter added to the model to improve fit destroys it.

Note also where this places us. `theta − beta` is a *difference on a logarithmic register*, which is the register CANON_INTERNAL §1 identifies as the one everything in this corpus lives in. The Rasch model was derived in the 1950s from a requirement about invariance, with no reference to surplus, pitch, or Fechner. **It arrived at the same register by an independent route, and §54.11 shows that the arrival is not a coincidence.**

54.0.3 Theorem 54.1 — item separability

The difference in log-odds between two items is independent of the person parameter.

```
ln[P_ni/(1−P_ni)] − ln[P_nj/(1−P_nj)]  =  (theta_n − beta_i) − (theta_n − beta_j)
                                       =  beta_j − beta_i
```

`theta_n` cancels identically. Not approximately — identically, for every person. Verified numerically with `beta_i = −0.80` and `beta_j = +0.60`:

theta	P_i	P_j	logit P_i	logit P_j	difference	odds ratio
-3.0	0.0998	0.0266	-2.2000	-3.6000	**+1.4000**	4.0552
-2.0	0.2315	0.0691	-1.2000	-2.6000	**+1.4000**	4.0552
-1.0	0.4502	0.1680	-0.2000	-1.6000	**+1.4000**	4.0552
0.0	0.6900	0.3543	+0.8000	-0.6000	**+1.4000**	4.0552
1.0	0.8581	0.5987	+1.8000	+0.4000	**+1.4000**	4.0552
2.0	0.9427	0.8022	+2.8000	+1.4000	**+1.4000**	4.0552
3.0	0.9781	0.9168	+3.8000	+2.4000	**+1.4000**	4.0552

The success probabilities range from 3% to 98%. **The comparison between the items is 1.4000 in every row.** Ask a struggling person or an expert; the items stand in the same relation.

54.0.4 Theorem 54.2 — person separability

The difference in log-odds between two persons is independent of the item parameter. By the same cancellation, with `theta_1 = 1.10` and `theta_2 = −0.30`:

item beta	logit P_1 - logit P_2
-2.0	**+1.4000**
-1.0	**+1.4000**

item beta	logit P_1 - logit P_2
0.0	**+1.4000**
+1.0	**+1.4000**
+2.0	**+1.4000**

`theta_1 − theta_2 = 1.4000`, whichever item you use. **This is the half of specific objectivity that matters for Volume Two: a person's measure does not depend on which items were used.**

54.0.5 The badge, and the methodological inversion

Note on status. The Rasch model is **a model**, and the theorems above are theorems *about the model*, not about people. They say: *if* responses follow this model, *then* separability holds. They say nothing about whether any real response set does follow it.

This produces the methodological inversion that defines Rasch measurement and separates it from mainstream item response theory:

```
ITEM RESPONSE THEORY:   choose the model that fits the data.
                        Add parameters until fit is acceptable.
                        Goal: describe the data well.

RASCH MEASUREMENT:      require the data to fit the model.
                        Where they do not, revise the INSTRUMENT.
                        Goal: obtain measurement, or find out you cannot.
```

Neither position is foolish, and the choice between them is a choice about what you are trying to do. If you want the best description of a data set, fit the richest model. If you want a *measuring instrument* — something whose readings transfer between contexts — you must insist on the model that has the invariance property, and treat misfit as a defect of the instrument rather than as a licence to elaborate the model. Chapter 53's Theorem 53.1 is the reason: the properties that make numbers usable are properties of the *scale*, and they are lost the moment the model allows item-by-person interaction.

This manual takes the Rasch position, and badges it as a position. The reason is stated plainly in §54.8 and demonstrated with numbers: models with free discrimination parameters fit better and do not measure.

54.0.6 The connection to Chapter 53

Chapter 53 §53.9 introduced additive conjoint measurement (Luce & Tukey 1964): a purely ordinal table on two crossed factors, satisfying double cancellation, licenses interval scales for both factors. The deterministic form is brittle — a single noisy response produces a violation.

The Rasch model is the probabilistic form of exactly that structure. Persons and items are the two crossed factors; the response table is the outcome; `theta_n − beta_i` is the additive combination; the logistic function supplies the error model that deterministic cancellation lacks. Fit statistics (§54.10) are the stochastic replacement for counting cancellation violations.

Rasch measurement is additive conjoint measurement with error, and that is why it delivers interval scales from ordinal data honestly rather than by fiat. It is the answer to Michell's objection, and it is the answer to a practitioner who has only ratings to work with.

The Engineering

54.1 Sufficiency — why separation is possible at all

Separability is not luck. It follows from a property of the model's likelihood, and the property is worth deriving because it is the mechanism.

For one person answering k items, with `xi = exp(theta)` and `eps_i = exp(−beta_i)`:

```
                    prod_i  (xi eps_i)^{x_i}          xi^r * prod_i eps_i^{x_i}
P(x_1..x_k | theta) = -------------------------  =  ----------------------
                       prod_i  (1 + xi eps_i)           prod_i (1 + xi eps_i)
```

where `r = sum_i x_i` is the RAW SCORE.

The person parameter enters only through `xi^r`. The response pattern enters only through `prod eps_i^{x_i}`. The likelihood factorises, and by the factorisation criterion **the raw score `r` is a sufficient statistic for `theta`.**

The consequence is the machinery of the whole chapter. Condition on `r`:

```
                              prod_i eps_i^{x_i}
P(x_1..x_k | r)  =  --------------------------
                              gamma_r(eps)
```

where `gamma_r` is the **elementary symmetric function** of order `r` in the `eps_i` — the sum of all products of `r` distinct `eps` values. **The person parameter has vanished completely.** The conditional distribution of the response pattern given the raw score depends on the item parameters alone.

Verified numerically. Five items with `beta = (−1.2, −0.5, 0.0, 0.6, 1.4)`, a person at `theta = 0.4`, four different response patterns all scoring `r = 3`:

pattern	r	P(pattern | theta)	P(pattern | r)
1 1 1 0 0	3	0.142349	0.375236
1 0 1 1 0	3	0.047384	0.124905
0 1 1 1 0	3	0.023530	0.062026
1 1 0 1 0	3	0.078123	0.205934

The middle column depends on `theta`. **The right-hand column does not — it is computable from the item parameters alone, and it is the quantity conditional estimation maximises.**

54.2 The gamma recursion

gamma_r looks expensive — it is a sum over all C(k, r) subsets — but there is a recursion that computes the whole vector in O(k^2). Add items one at a time; each new item either contributes to the product or does not:

```
gamma_r^(m)  =  gamma_r^(m-1)  +  eps_m * gamma_{r-1}^(m-1)
```

Worked in full for the five items above, eps = (3.32012, 1.64872, 1.00000, 0.54881, 0.24660):

```
after item 1:  1.00000  3.32012
after item 2:  1.00000  4.96884   5.47395
after item 3:  1.00000  5.96884  10.44279   5.47395
after item 4:  1.00000  6.51765  13.71855  11.20507  3.00417
after item 5:  1.00000  6.76425  15.32579  14.58802  5.76730  0.74082

check:  gamma_1 = sum(eps)  = 6.76425     OK
        gamma_5 = prod(eps) = 0.74082     OK
```

The two checks at the bottom are the ones to run in any implementation: the first-order symmetric function is the sum, the top-order one is the product. **Numerical instability in these functions is the commonest bug in Rasch software and the two checks catch most of it.**

54.3 Conditional maximum likelihood

Maximise the product of conditional probabilities over persons:

```
                  N   [    prod_i eps_i^{x_ni}   ]
ln L_C  =  ln  prod    [ ------------------------ ]
                 n=1   [      gamma_{r_n}(eps)   ]

        =  - sum_i  s_i beta_i  -  sum_n ln gamma_{r_n}(eps)
```

where s_i = sum_n x_ni is the ITEM SCORE.

Three properties of this estimator matter and should be stated:

CML-1 No person parameters appear. There are N of them and they are ALL eliminated. This is why CML is consistent as N grows -- unlike joint estimation, where the number of nuisance parameters grows with the sample and the estimator is biased.

CML-2 The item score s_i is a sufficient statistic for beta_i. Symmetric to the person side, as specific objectivity requires.

CML-3 Persons with extreme scores (r = 0 or r = k) contribute nothing: gamma_0 and gamma_k are constants and the conditional likelihood is degenerate. They are EXCLUDED, and the exclusion must be reported.

The atom of the whole procedure is visible in the two-item case. Among persons who got exactly one of items i and j right:

$$P(\text{item } i \text{ right} \mid r = 1) = \frac{\exp(-\beta_i)}{\exp(-\beta_i) + \exp(-\beta_j)} = 0.802184$$

with $\beta_i = -0.80$, $\beta_j = +0.60$. Log odds $= \beta_j - \beta_i = 1.4000$.

No theta appears anywhere in that expression. The relative difficulty of two items is estimated from the persons who split them, and how able those persons are does not enter. That is specific objectivity as an estimation procedure rather than as a theorem.

54.4 Calibration — a worked recovery

Twelve items with true difficulties evenly spaced from −2.00 to +2.00 logits (centred), 2000 simulated persons drawn Normal(0, 1.2), 1945 usable after excluding extreme scores. CML estimates:

item	true beta	CML beta	error
1	-2.0000	-1.8690	+0.1310
2	-1.6364	-1.5544	+0.0820
3	-1.2727	-1.3236	-0.0509
4	-0.9091	-0.9090	+0.0001
5	-0.5455	-0.6309	-0.0854
6	-0.1818	-0.1432	+0.0386
7	+0.1818	+0.1336	-0.0482
8	+0.5455	+0.5250	-0.0204
9	+0.9091	+0.9072	-0.0019
10	+1.2727	+1.4011	+0.1284
11	+1.6364	+1.5664	-0.0699
12	+2.0000	+1.8967	-0.1033

RMS error = 0.0764 logits max |error| = 0.1310 logits
correlation(true, estimated) = 0.998382

54.5 Separation and reliability — how much measurement you actually have

Estimation accuracy is not the same as measurement precision, and Rasch practice reports both.

$$\text{person } SE_n = 1 / \sqrt{\sum_i P_{ni}(1 - P_{ni})}$$
$$\text{item } SE_i = 1 / \sqrt{\sum_n P_{ni}(1 - P_{ni})}$$

$$\text{true SD} = \sqrt{\text{observed SD}^2 - \text{mean } SE^2}$$
$$\text{SEPARATION index } G = \text{true SD} / \text{RMSE}$$
$$\text{RELIABILITY} = G^2 / (1 + G^2)$$
$$\text{distinct STRATA} = (4G + 1) / 3$$

For the calibration above:

observed SD of person measures = 1.3675 logits
RMSE of person measures = 0.7576 logits

```
true SD                        = 1.1385 logits
person separation G            = 1.5028
person reliability             = 0.6931
distinct strata                = 2.34

item separation                = 22.0854
item reliability               = 0.9980
mean item SE                   = 0.0578 logits
```

Read the last line of the person block: 2.34 strata. Twelve items, two thousand people, and the instrument distinguishes roughly two and a third levels of the attribute. It can tell high from low and not much else. The item side is superb — 0.0578 logits of item error — because two thousand people calibrated twelve items. **Item precision and person precision are bought by different resources: items are calibrated by persons, persons are measured by items.**

This is the number practitioners most often fail to report, and it is the number that determines whether the instrument can do the job. **An instrument with a person reliability of 0.6931 cannot support individual decisions.** It can support group comparisons. Saying which is a professional obligation.

54.6 The invariance test — running FVP-7 on an instrument

The theorems in §54.0.3 and §54.0.4 are about the model. The test is about the data. **Split the persons into disjoint groups that differ substantially on the attribute, calibrate the items separately in each, and compare.**

The person sample above was split at the median ability (lower half mean `theta = −0.949`, upper half mean `theta = +0.872` — a separation of 1.82 logits, so the two groups are genuinely different populations):

item	`beta` (low sample)	`beta` (high sample)	difference
1	-1.8684	-1.8714	-0.0030
2	-1.4895	-1.7360	-0.2465
3	-1.3595	-1.2439	+0.1156
4	-0.9382	-0.8568	+0.0814
5	-0.7165	-0.5056	+0.2109
6	-0.1480	-0.1363	+0.0118
7	+0.1505	+0.1203	-0.0302
8	+0.5995	+0.4763	-0.1232
9	+1.0241	+0.8481	-0.1760
10	+1.4439	+1.3856	-0.0582
11	+1.5202	+1.5849	+0.0647
12	+1.7820	+1.9349	+0.1529

```
mean difference = +0.0000   sd of difference = 0.1360   max |difference| = 0.2465
correlation(low, high) = 0.994399
```

Two disjoint populations, differing by 1.82 logits in ability, produced item calibrations correlating at 0.9944 with a mean difference of exactly zero and no difference exceeding a quarter

of a logit. That is what specific objectivity looks like in data. **This single table is the strongest form of evidence available in psychometrics, and it is the exact analogue of FVP-7.**

The formal statistic is a per-item invariance `t`:

```
t_i = ( beta_i^{high} − beta_i^{low} ) / sqrt( SE_i^{high 2} + SE_i^{low 2} )
```

Flag at `|t| > 2.58` (roughly a 1% two-sided criterion, appropriate under multiple testing across a short instrument). §54.12 runs it on data with planted defects.

54.7 What breaks separability — the two-parameter model

The two-parameter logistic model (Birnbaum; Lord & Novick, *Statistical Theories of Mental Test Scores*, 1968) adds a discrimination parameter per item:

```
P(X_ni = 1) = exp( a_i (theta_n − beta_i) ) / [ 1 + exp( a_i (theta_n − beta_i) ) ]
```

It fits real data better. It fits real data better because it has more parameters. **And it destroys specific objectivity outright**, which is not a subtle statistical cost but a total loss of the property that made the numbers measurements.

The item log-odds difference under 2PL:

```
a_i (theta − beta_i) − a_j (theta − beta_j) = (a_i − a_j) theta + (a_j beta_j − a_i beta_i)
```

The `theta` term does not cancel unless `a_i = a_j`. The same two items as §54.0.3, now with `a_i = 1.8` and `a_j = 0.6`:

theta	P_i	P_j	logit P_i	logit P_j	difference	odds ratio
-3.0	0.0187	0.1034	-3.9600	-2.1600	**-1.8000**	0.1653
-2.0	0.1034	0.1736	-2.1600	-1.5600	**-0.6000**	0.5488
-1.0	0.4110	0.2769	-0.3600	-0.9600	**+0.6000**	1.8221
0.0	0.8085	0.4110	+1.4400	-0.3600	**+1.8000**	6.0496
1.0	0.9623	0.5597	+3.2400	+0.2400	**+3.0000**	20.0855
2.0	0.9936	0.6985	+5.0400	+0.8400	**+4.2000**	66.6863
3.0	0.9989	0.8085	+6.8400	+1.4400	**+5.4000**	221.4064

The difference sweeps from −1.80 to +5.40 and **changes sign at**

```
theta* = (a_i beta_i − a_j beta_j) / (a_i − a_j) = −1.5000
```

Below `theta = −1.5`, item `i` is harder. Above it, item `j` is harder. Which item is more difficult is not a property of the items. It depends on who you ask. The odds ratio moves by a factor of over 1300 across the range.

This is the whole argument in one table. **A model that fits better and cannot answer "which item is harder" without naming a person has not measured the items.** The 2PL is a perfectly good descriptive model and a perfectly bad measuring instrument, and Volume Two needs a measuring instrument.

54.8 The 2PL under the invariance test

The theoretical point becomes an empirical diagnostic. Data were generated from a 2PL — items 1-6 with a = 1.8, items 7-12 with a = 0.6, same difficulties as before — and then put through exactly the Rasch calibration and split-sample procedure of §54.6:

item	beta (low)	beta (high)	difference	true a
1	-2.2508	-3.5274	-1.2766	1.8
2	-1.6428	-2.3178	-0.6750	1.8
3	-1.1397	-1.9441	-0.8043	1.8
4	-0.5026	-1.3115	-0.8089	1.8
5	+0.0621	-0.6818	-0.7438	1.8
6	+0.7035	-0.0222	-0.7257	1.8
7	+0.2261	+0.9978	+0.7717	0.6
8	+0.4828	+1.3075	+0.8248	0.6
9	+0.7755	+1.4943	+0.7189	0.6
10	+0.8497	+1.8437	+0.9940	0.6
11	+1.1991	+2.0240	+0.8249	0.6
12	+1.2371	+2.1374	+0.9002	0.6

```
mean |difference|  =  0.8391   (Rasch data: 0.1062)   ->  7.9 times larger
max  |difference|  =  1.2766   (Rasch data: 0.2465)
```

Look at the sign column. Every high-discrimination item drifted negative. Every low-discrimination item drifted positive. Not one exception in twelve. **The drift is not noise; it is a signature, and the signature reads out the discrimination structure directly.** The test does not merely detect the violation, it tells you what kind of violation it is.

That is the practical value of the invariance test and the reason it belongs in every FVP-7 report. A global fit index would have said "acceptable". The split-sample table says "your items have heterogeneous discriminations, the high ones are items 1 through 6, and your scale is not a scale."

54.9 Fit statistics — what misfit tells you

The residual of a single response, standardised:

```
z_ni = ( x_ni − P_ni ) / sqrt( P_ni (1 − P_ni) )    mean 0, variance 1 under the model
```

Two summaries, per item:

```
                 1
OUTFIT MNSQ  =  ---  sum_n  z_ni^2                    UNWEIGHTED mean square
                 N

                 sum_n ( x_ni − P_ni )^2
INFIT  MNSQ    =  -------------------------                  INFORMATION-
WEIGHTED mean square
                 sum_n  P_ni (1 − P_ni)
```

Both have expectation 1 under the model. The difference between them is which responses dominate:

```
OUTFIT is OUTLIER-SENSITIVE. Its denominator P(1-P) is tiny for far-off-
target responses,
     so a single unexpected response from a badly mismatched person can move it a lot.
       Use it to detect: guessing, careless errors, miskeys, lucky hits.

INFIT is INLIER-SENSITIVE. It weights each residual by the information that response
       carries, so on-target responses dominate.
     Use it to detect: problems with the items people are actually well matched to -
-
       the failures that matter for the measure.
```

Interpretation of the mean square, with the two directions carrying entirely different diagnoses:

```
MNSQ >> 1     UNDERFIT -- more randomness than the model predicts.
         Off-construct item, ambiguous wording, guessing, careless responding,
         multidimensionality. THIS IS THE DANGEROUS DIRECTION: it degrades measurement.

MNSQ << 1     OVERFIT -- LESS randomness than the model predicts.
           Responses too predictable: item dependency, redundancy, a Guttman-
like
         over-discriminating item. Does not degrade measurement, but it INFLATES
             reliability and makes the instrument look better than it is.
```

The asymmetry is the practically important part and it is routinely inverted by newcomers, who treat MNSQ = 0.4 as a triumph. It is not. It means an item is telling you something you already knew, and it means your reported reliability is a fiction.

Standardised versions (ZSTD) convert the mean square to an approximate normal deviate via the Wilson-Hilferty cube-root transformation. Report both: **the mean square says how much misfit, the ZSTD says how sure you are it is real.** With large samples, trivial misfit becomes highly significant, so ZSTD alone will condemn everything.

54.10 Fit statistics on data with planted defects

Twelve items, 2958 usable persons, nine clean Rasch items and three deliberately broken ones.

item	beta	OUTFIT MNSQ	ZSTD	INFIT MNSQ	ZSTD	pt-measure r	planted defect
1	-1.851	0.859	-1.92	0.911	-1.96	+0.419	clean
2	-1.571	0.897	-1.57	0.886	-2.68	+0.465	clean
3	-1.291	0.925	-1.26	0.939	-1.47	+0.465	clean
4	-1.522	0.936	-0.98	1.000	+0.00	+0.404	**guessing, c = 0.30**
5	-0.587	0.918	-1.76	0.944	-1.45	+0.519	clean
6	-0.210	0.953	-1.06	0.936	-1.70	+0.535	clean
7	+0.168	0.887	-2.64	0.915	-2.29	+0.556	clean

item	beta	OUTFIT MNSQ	ZSTD	INFIT MNSQ	ZSTD	pt-measure r	planted defect
8	+0.158	**1.411**	**+8.23**	**1.252**	**+6.14**	+0.328	**low discrimination, a = 0.35**
9	+0.896	0.913	-1.74	0.947	-1.31	+0.510	clean
10	+1.302	0.850	-2.68	0.924	-1.80	+0.496	clean
11	+2.599	**0.404**	**-7.45**	**0.683**	**-6.09**	+0.487	**overdiscriminating, a = 3.0**
12	+1.909	0.788	-2.99	0.894	-2.24	+0.451	clean

Item 8 is caught cleanly: `OUTFIT = 1.411`, `ZSTD = +8.23`, and the point-measure correlation drops to `+0.328` against a clean-item range of `0.40` to `0.56`. Classic underfit — the item is noisy relative to the construct.

Item 11 is caught in the other direction: `OUTFIT = 0.404`, `ZSTD = −7.45`. **Severe overfit. Its responses are far too predictable, and if you read low mean squares as good news you would have called it your best item.** Note also its estimated difficulty, `+2.599` — well outside the true range — because forcing a slope-3 item into a slope-1 model displaces its location.

Item 4 is the finding of the section, and it is a negative one. The guessing item, with a 30% lower asymptote, posted `OUTFIT = 0.936` and `INFIT = 1.000`. **Both statistics say "perfect".** By fit statistics alone, item 4 is the best-behaved item on the instrument. It is broken.

No single fit statistic is sufficient, and the invariance test is the criterion that fit statistics only approximate. §54.12 runs the invariance `t` on this same data set and catches item 4 at `t = +4.14`. That is why FVP-7 is a separate gate rather than a footnote to model fit, and it is why Chapter 52 insisted it is the gate that decides.

An honest caveat about the mean squares. The clean items in the table average `OUTFIT = 0.888`, not `1.000`. This is not evidence that they misfit. It is the deflation caused by estimating both `theta` and `beta` from the same responses: fitted parameters absorb some residual variance. The same data scored against the *true* parameters gives:

```
mean OUTFIT with true theta and true beta = 0.9995     (expectation 1.000)
mean INFIT  with true theta and true beta = 0.9988
```

The expectation of 1.000 is exact for the model, and the shortfall to 0.888 is an artefact of estimation, not a property of the items. Anyone comparing observed mean squares to a nominal 1.000 without accounting for this will systematically over-diagnose overfit. Compare items to the *instrument's own* mean, and to each other.

54.11 Why this is exactly what Volume Two needs

Now the payoff, and it is a structural result rather than an analogy.

Step 1. Rasch measures live on an interval scale in the log-odds register. The model `ln[P/(1−P)] = theta − beta` determines `theta` up to an additive constant: the origin is arbitrary (fixed by convention, usually by centring the items at zero), and the *unit* is fixed by the model, because the discrimination is constrained to 1. Choosing a different logarithm base rescales all measures by one common factor — exactly the way nats and bits differ (CANON_INTERNAL §1) — and is a change of unit, not a change of scale.

Step 2. Identify the quantity with the exponential of the measure. *Postulate SO1, badged as a postulate:* if `D` and `P` are calibrated as Rasch measures `theta_D` and `theta_P` on one common scale, then the corresponding magnitudes are `D = exp(theta_D + b)` and `P = exp(theta_P + b)`, where `b` is the scale's arbitrary common origin. This identification is a modelling choice. It is not derived, and it is badged here.

Step 3. The arbitrary origin cancels from the ratio. Verified across four wildly different origin choices, with `theta_D = 1.45` and `theta_P = 0.62`:

origin b	D	P	D/P	Delta = |ln r|
+0.0	4.2631	1.8589	**2.293319**	**0.830000**
-3.0	2.1225e-01	9.2551e-02	**2.293319**	**0.830000**
+2.5	5.1935e+01	2.2646e+01	**2.293319**	**0.830000**
+100.0	1.1460e+44	4.9970e+43	**2.293319**	**0.830000**

```
D/P  =  exp(theta_D − theta_P)  =  exp(0.8300)  =  2.293319
```

The magnitudes `D` and `P` range over forty-four orders of magnitude and the ratio does not move. The arbitrary constant that makes the logit scale merely *interval* is exactly the constant that cancels in a ratio.

Step 4. The distortion measure is on a genuine ratio scale.

```
Delta  =  | ln(D/P) |  =  | theta_D − theta_P |
```

Its true zero is not a convention: `Delta = 0` if and only if `theta_D = theta_P`, which is the calibrated state, and that identity holds in every unit and under every origin. The only remaining freedom is the unit:

```
Delta in logits (nats) = 0.830000
Delta in bits          = 1.197437        (divide by ln 2 = 0.693147)
ratio                  = 0.693147        a pure multiplicative constant
```

A quantity with a fixed true zero and a free multiplicative unit is a ratio scale, by Chapter 53's definition. Therefore:

> **Theorem 54.3 (given Postulate SO1).** *If `D` and `P` are measured as Rasch person-measures on one common scale, then `Delta = |ln(D/P)| = |theta_D − theta_P|` is ratio-scaled with a true zero at perfect calibration, and `D/P` is determined up to the choice of unit.*

This is the single most important result in Part X for this manual's own formulas. Chapter 53 showed that ordinal data cannot support D/P. Chapter 54 shows there is a constructive route from ordinal responses to a properly scaled distortion measure: build a Rasch instrument, calibrate D and P on a *common* scale, and read the distortion off as a logit difference. **The raw ordinal responses never became ratio data. The measure derived from them did, because the model earned it.**

And note where this lands. Distortion is a *difference of logarithms* — which is what CANON_INTERNAL §1 said every quantity in this corpus is, and what Chapter 36's Lens Theorem said perception operates in. **Rasch arrived at the logarithmic register in 1960 from a pure invariance requirement about test items, with no reference to Fechner, to surplus, or to anything in this book.** Two independent derivations, one register. That convergence is worth stating and worth not overstating: it is a structural correspondence between formalisms, exactly like the predictive-processing correspondence of CANON_INTERNAL §4, and structural similarity is not identity.

54.12 What it would cost — the precision arithmetic

The result above is constructive, which means it can be costed. Person measurement error in the best case (every item targeted at the person, P = 0.5, information 0.25 per item):

```
SE(theta)  =  1 / sqrt( k * 0.25 )  =  2 / sqrt(k)
SE(Delta)  =  SE(theta_D) * sqrt(2)          (two independent measures)
```

items per term k	best-case SE(theta)	SE(Delta)	smallest detectable Delta (95%)
5	0.8944	1.2649	2.4792
10	0.6325	0.8944	1.7531
12	0.5774	0.8165	1.6003
20	0.4472	0.6325	1.2396
30	0.3651	0.5164	1.0121
44	0.3015	0.4264	0.8357
60	0.2582	0.3651	0.7157
100	0.2000	0.2828	0.5544
150	0.1633	0.2309	0.4526
200	0.1414	0.2000	0.3920

Inverted, so it can be read as a design requirement:

distortion to detect	ratio r	items required per term
1.00 logits	2.718	31
0.50 logits	1.649	123
0.30 logits	1.350	342
0.20 logits	1.221	769
0.10 logits	1.105	3074

This table is sobering and it is meant to be. A twenty-item instrument can only detect distortions above about `1.24` logits — a frame-to-intake ratio of about 3.5. Anything subtler is inside the noise. To resolve a ratio of `1.35` you need three hundred and forty-two items per term, which no one will sit through in one session and which therefore implies adaptive testing, item banking, or repeated measurement across occasions.

A practitioner reporting `r = 1.2` from a twelve-item pair of scales is reporting noise, and the arithmetic above says so exactly. That is the kind of statement the Formula Validation Protocol exists to make available before the study is run rather than after.

The Specification

SO-1 — State the model and its constraint. Report that the discrimination is fixed at 1 and that this is a deliberate measurement choice, not an oversight. Anyone using a 2PL or 3PL must state that specific objectivity has been abandoned and say why.

SO-2 — Estimate by conditional maximum likelihood where feasible. CML eliminates the person parameters and is consistent as `N` grows. Where marginal or joint methods are used instead, report the choice and its bias implications.

SO-3 — Report the gamma checks. `gamma_1 = sum(eps)` and `gamma_k = prod(eps)`. Numerical failure in the elementary symmetric functions is the commonest implementation bug.

SO-4 — Exclude and report extreme scores. Persons at `r = 0` or `r = k` and items answered correctly by everyone or no one contribute nothing to CML. State how many were excluded. A large exclusion count means the instrument is mistargeted.

SO-5 — Report person AND item separation, reliability, and strata. Not reliability alone. `(4G+1)/3` strata is the number that tells a reader what the instrument can actually distinguish. An instrument reported without its strata count is not specified.

SO-6 — Never claim individual-level decisions below person reliability 0.90. The 0.69 example in §54.5 supports group comparison and nothing finer.

SO-7 — Run the split-sample invariance test, always. Split on the attribute itself, not on convenience. Report the full item-by-item table, the mean difference, the sd of differences, the maximum, and the correlation. **This is the FVP-7 deliverable and it is not optional.**

SO-8 — Split on a second variable too. Attribute-based splitting detects discrimination heterogeneity. Splitting on group membership detects differential item functioning. Both must be run and both reported.

SO-9 — Report the invariance `t` per item. `t_i = (beta_i^A − beta_i^B) / sqrt(SE_A^2 + SE_B^2)`, with the flagging threshold stated in advance and justified against the number of items tested.

SO-10 — Report OUTFIT and INFIT, mean square and standardised, for every item. Four numbers per item. Not a global fit index, not a selection of the good ones.

SO-11 — Interpret overfit as a defect. `MNSQ` well below 1 means dependency or redundancy and it inflates reliability. Treat it as a finding, never as a success.

SO-12 — Correct for estimation deflation before diagnosing. Compare items to the instrument's own mean mean-square, not to the nominal 1.000. §54.10 measured the artefact at roughly `0.11`.

SO-13 — Report the point-measure correlation. It catches things the mean squares miss and it is free.

SO-14 — Do not rely on fit statistics alone. Item 4 of §54.10 passed both mean squares and was broken. Fit statistics are a screen; invariance is the criterion.

SO-15 — Test unidimensionality explicitly. Principal-components analysis of the standardised residuals, with the variance explained by the first residual contrast reported. Rasch fit does not by itself establish that one dimension is present.

SO-16 — Calibrate `D` and `P` on ONE common scale. Either one item bank measuring both, or two banks linked by common items or common persons, with the linking constants and their errors reported. **Two independently centred scales cannot be differenced**, because each has its own arbitrary origin and the origins do not cancel. This is Chapter 53's S-5 in Rasch form and it is the single most likely point of failure in an applied study.

SO-17 — Report `Delta` in logits with its standard error. `SE(Delta) = sqrt(SE_D^2 + SE_P^2)`. A distortion is not reportable unless its interval excludes zero.

SO-18 — Use the precision table before the study, not after. §54.12. Decide what magnitude of distortion matters, read off the item count, and either build that instrument or narrow the claim.

SO-19 — Badge Postulate SO1 wherever the exponential identification is used. The step from a Rasch logit measure to a magnitude `D = exp(theta + b)` is a modelling choice and carries its badge under CANON §2.

SO-20 — State the unit. Logits (nats) or bits. `Delta` is unit-dependent by a multiplicative constant and a reported number without its unit is not a number.

The Failure Modes

FM-54.1 . THE 2PL ESCAPE

Damages: specific objectivity, totally, in exchange for a better fit index.

Items misfit. A discrimination parameter is added per item. Fit improves. The report says the model now describes the data well, and it does — and the scale is gone.

Diagnostic signature: the model was chosen after inspecting fit; the discrimination parameters vary substantially across items; no invariance test is reported; the phrase "better model fit" appears without the phrase "loss of parameter separation".

Detecting measurement: §54.7's table on your own items. Compute the item log-odds difference at `theta = −2` and at `theta = +2`. If it moves, item difficulty is person-dependent.

Fix: revise the instrument, not the model. Remove or rewrite the items whose discriminations deviate. **If you genuinely need the 2PL to describe your data, say that you have a description and not a measure, and stop reporting the numbers as if they transferred.**

FM-54.2 . THE UNTESTED INVARIANCE

Damages: FVP-7 and SO-7.

A Rasch analysis is run, fit statistics are reported, and no split-sample calibration is ever performed. The theorem is cited in the introduction; its empirical test is absent from the results.

Diagnostic signature: the words "specific objectivity" appear in the introduction and nowhere in the methods. No table compares calibrations across samples.

Detecting measurement: look for the table. There either is one or there is not.

Fix: run it. It costs one extra estimation and it is the most informative table in the report. §54.6 and §54.8 show what a pass and a failure look like side by side, and they are not subtle.

FM-54.3 . THE FIT-STATISTIC ALIBI

Damages: SO-14, and the credibility of everything downstream.

Acceptable mean squares are treated as proof of measurement. They are not. §54.10's item 4 — a guessing item with a 30% lower asymptote — returned `OUTFIT = 0.936` and `INFIT = 1.000`, the best numbers on the instrument.

Diagnostic signature: the conclusion "the data fit the Rasch model" rests entirely on mean squares in a conventional range; no invariance test; no residual dimensionality analysis.

Detecting measurement: the invariance `t` (SO-9). On §54.10's data it flags item 4 at `t = +4.14` where both mean squares said nothing.

Fix: SO-7 through SO-15. Three independent checks — fit, invariance, dimensionality — and no one of them substitutes for another.

FM-54.4 . READING OVERFIT AS SUCCESS

Damages: reported reliability, and the instrument's real precision.

An item posts `MNSQ = 0.40` and is celebrated. It is over-discriminating or dependent on another item. Reliability is inflated because the responses are more predictable than the construct warrants.

Diagnostic signature: mean squares clustered well below 1; unusually high reliability for a short instrument; items that are near-paraphrases of one another.

Detecting measurement: residual correlations between item pairs. Strong positive residual correlation is local dependence and it is the usual cause.

Fix: combine or remove the dependent items and recalibrate. Expect reliability to fall; the fall is the correction, not the damage.

FM-54.5 . THE UNLINKED SCALES

Damages: SO-16, and with it every `Delta` the study reports.

`D` and `P` are calibrated by two separate Rasch analyses, each centred on its own items at zero. The difference `theta_D − theta_P` is then computed as though the origins were shared. They are not, and each was set arbitrarily.

Diagnostic signature: two calibrations, each reported with mean item difficulty exactly zero, and no linking design described. The mean `Delta` across a population comes out implausibly close to zero, or implausibly far from it, with no substantive explanation.

Detecting measurement: look for the linking constant. If there is no common item set and no common person set, there is no link.

Fix: a linking design — common items across both banks, or a common person sample taking both — with the linking constant and its standard error reported. **Without a link, `Delta` is the difference of two arbitrary constants and means nothing.** This is the same defect as Chapter 53's FM-53.6, arriving through a different door.

FM-54.6 . THE MISTARGETED INSTRUMENT

Damages: person precision, silently.

The items sit at difficulties far from where the persons sit. Everyone scores near the top or near the bottom. Extreme scores are excluded, the remaining sample is small and unrepresentative, and person standard errors are enormous.

Diagnostic signature: a large excluded fraction; a person-item map with the two distributions barely overlapping; person separation below 1.5 and strata below 2.

Detecting measurement: SO-4 and SO-5. The exclusion count and the strata count together are diagnostic.

Fix: add items at the person distribution's location. **Item precision is bought with persons and person precision is bought with items; a mistargeted instrument buys neither.**

FM-54.7 . CENTRING CONTAGION

Damages: every item estimate, by a common shift.

The scale's origin is set by centring the item difficulties. One badly misfitting item drags the mean, and therefore shifts every other item's estimate. Clean items are then flagged as misbehaving.

Diagnostic signature: several items flagged at moderate `t` alongside one flagged at extreme `t`; the moderate flags disappear when the extreme item is removed and the scale recentred.

Detecting measurement: remove the worst-fitting item, recalibrate, and re-run the invariance test. In §54.12's run, item 5 was flagged at `t = +3.71` in the presence of three planted defects; the flag is contamination, not a property of item 5.

Fix: anchor the origin on a stable subset of items rather than on the full-set mean, and always re-run diagnostics after removing an item. **Never remove several items at once: each removal moves the scale.**

FM-54.8 . THE ASSUMED DIMENSION

Damages: the meaning of the measure, entirely.

Fit statistics are acceptable, invariance holds, and the instrument is nonetheless measuring two things at once whose mixture happens to be stable in the sample studied. The measure will not transfer to a population where the mixture differs.

Diagnostic signature: residual principal components with a first contrast explaining substantial variance; items clustering by content in the residual space; measures that behave differently in a new population despite fitting in the old one.

Detecting measurement: SO-15. Principal-components analysis of the standardised residuals, with the first-contrast eigenvalue reported.

Fix: split the instrument along the contrast and calibrate each dimension separately. Two honest measures beat one composite whose composition is unknown.

Worked Example

The full sequence on the twelve-item instrument with planted defects, ending in the FVP-7 verdict and the application to `D/P`.

54.13.1 Calibration and fit

The instrument, the fit table, and the deflation correction are §54.10. Summary of what the fit screen alone concluded:

```
UNDERFIT flagged:   item 8    OUTFIT 1.411  ZSTD +8.23   pt–measure +0.328
OVERFIT  flagged:   item 11   OUTFIT 0.404  ZSTD –7.45
CLEAN by fit:     all ten others, INCLUDING item 4 (OUTFIT 0.936, INFIT 1.000)
```

Two of three planted defects detected. **The third — the guessing item — passed cleanly.**

54.13.2 The invariance test on the same data

Persons split at the median measure, 1479 in each half, items recalibrated independently in each, with standard errors:

item	b_low	SE	b_high	SE	difference	t	flag
1	-2.056	0.061	-1.972	0.101	+0.084	+0.71	
2	-1.791	0.059	-1.630	0.089	+0.161	+1.51	
3	-1.478	0.057	-1.417	0.083	+0.061	+0.60	
4	-1.816	0.059	-1.395	0.083	+0.421	**+4.14**	**DIF**
5	-0.857	0.057	-0.532	0.066	+0.324	**+3.71**	**DIF**
6	-0.392	0.059	-0.261	0.063	+0.131	+1.51	
7	-0.013	0.063	+0.136	0.060	+0.149	+1.72	
8	-0.645	0.058	+0.705	0.057	+1.350	**+16.55**	**DIF**
9	+0.819	0.076	+0.830	0.057	+0.012	+0.12	

item	b_low	SE	b_high	SE	difference	t	flag
10	+1.129	0.084	+1.292	0.058	+0.163	+1.60	
11	+5.212	0.505	+2.384	0.068	-2.828	**-5.55**	**DIF**
12	+1.887	0.109	+1.860	0.062	-0.028	-0.22	

Flagged at `|t| > 2.58`: **items 4, 5, 8, 11.**

Read this against the planted defects:

```
item 4   guessing, c = 0.30      CAUGHT by invariance (t = +4.14), MISSED by fit
item 8    low discrimination        CAUGHT by both (t = +16.55, OUTFIT 1.411)
item 11  over-discriminating        CAUGHT by both (t = -5.55, OUTFIT 0.404)
item 5   clean                      FALSE FLAG -- centring contagion, FM-54.7
```

The invariance test caught all three planted defects. The fit statistics caught two. That is the empirical basis for SO-14 and for Chapter 52's insistence that FVP-7 is the deciding gate.

Item 5's false flag is instructive and is reported rather than tidied away. Three misfitting items distorted the centring, shifting every estimate, and the shift showed up as apparent DIF in a clean item. **The correct procedure is FM-54.7's fix: remove the worst item, recentre, re-run, and repeat one item at a time.** Removing all four flagged items simultaneously would have discarded a good item on the strength of a contamination artefact.

Item 11's `b_low = +5.212` with `SE = 0.505` is the over-discriminating item behaving exactly as theory predicts: for the low-ability half, a slope-3 item is nearly impossible, so its estimated difficulty runs away and its error explodes. **An item estimate with an error an order of magnitude larger than its neighbours is itself a diagnostic and should never be reported without comment.**

54.13.3 The verdict on the instrument

```
FIT SCREEN            items 8 and 11 fail
INVARIANCE TEST       items 4, 8, 11 fail; item 5 flagged as contagion
DIMENSIONALITY       not run in this demonstration (SO-15 would be required)

INSTRUMENT VERDICT   NOT MEASURING as constituted. Items 4, 8 and 11 must be
                rewritten or removed, one at a time with recalibration between,
                    and the invariance test re-run on the revised instrument.
```

Three broken items out of twelve is an ordinary result for a first draft of an instrument, and finding them is what the procedure is for. The failure here is not a failure of the method; it is the method working.

54.13.4 Applying it to `D/P`

Suppose the revised instrument survives. Two Rasch banks — one for perspective, one for perception — **linked by a common item set so that the two scales share one origin (SO-16).** A respondent is measured on both.

```
theta_D = +1.45 logits    theta_P = +0.62 logits    (common scale)

Delta   = theta_D - theta_P  =  0.8300 logits
D/P     = exp(0.8300)        =  2.293319
verdict = OVER-FRAMED
```

The origin invariance was verified in §54.11: shifting the common origin by −3.0, +2.5 or +100.0 leaves D/P = 2.293319 and Delta = 0.830000 unchanged to six decimals, while the implied magnitudes range across forty-four orders of magnitude. **The reported quantity is a property of the person, not of the scale's arbitrary centre.**

Now the precision, which is where honesty is required. With person standard errors of 0.30 logits each:

```
SE(Delta) = sqrt(0.30^2 + 0.30^2) = 0.4243 logits
95% CI    = 0.8300 +/- 1.96 * 0.4243 = [-0.0016, 1.6616]
```

The interval includes zero, barely. A distortion of 0.83 logits — a frame-to-intake ratio of 2.29, which is substantial over-framing — **is not distinguishable from perfect calibration at this precision.** From §54.12's table, standard errors of 0.30 logits correspond to roughly 44 targeted items per term.

The correct report, per SO-17 and SO-20, is therefore:

```
Delta = 0.83 logits, 95% CI [-0.00, 1.66], SE 0.42
   [Rasch person measures, common-linked banks, 44 items per term,
    person reliability to be reported, unit = logits (nats)]
   Interval includes zero: over-framing is INDICATED, NOT ESTABLISHED.
   To establish a distortion of this size at 95%, the precision table (54.12)
   requires approximately 123 items per term.
```

That report is worth more than a confident "r = 2.29" would be, and it is worth more precisely because it declines to claim what it cannot support.

54.13.5 What this chapter establishes, and what it does not

Established:

1. Specific objectivity is a real, checkable property, and the Rasch model has it exactly (Theorems 54.1, 54.2).
2. It is testable in data, and the split-sample invariance table is the test. Rasch data: mean difference 0.0000, sd 0.1360, correlation 0.9944. 2PL data: mean |difference| 7.9 times larger, with a systematic sign pattern.
3. Fit statistics are a screen, not a criterion. Three planted defects: fit caught two, invariance caught three.
4. Given Postulate SO1, Delta = |theta_D − theta_P| is RATIO-SCALED with a true zero at calibration, and D/P is determined up to the choice of unit (Theorem 54.3). This is a constructive route to the scale Chapter 53 demanded.
5. The precision required is large and computable: 31 items per term to detect a 1-logit distortion, 123 for half a logit, 342 for 0.3.

Not established:

1. That any real instrument for perspective or perception fits the Rasch model. Nobody has built one. FVP-7 for D and P: NOT YET RUN.
2. That perspective and perception are unidimensional. This must be tested, not assumed (SO-15), and it is the assumption most likely to fail.
3. That Postulate SO1's exponential identification is correct. It is a modelling choice, badged, and an alternative identification would change the scale type of D/P.
4. Anything at all about the surplus law, the storage law, or Paper A. Those are proved mathematics (CANON §1) and stand independently of this chapter.

The distance between the established list and the unestablished one is the honest measure of Volume Two's current scientific standing, and it is a distance measured in laboratory work rather than in argument. Chapter 55 specifies the falsification suite and Chapter 56 costs the programme.

What Part X has shown is that the route exists. **Psychology and sociology are not pseudo-science; they are sciences with a measurement problem and a replication problem, and the problem is tractable.** Psychophysics built ratio scales for sensation and they replicate. Rasch built parameter separation for attainment and it can be tested in an afternoon on any data set. Neither achievement required abandoning the study of subjective experience. **Both required taking the measurement seriously enough to let it fail.**

That is the whole of the discipline this manual asks for, and it is the reason these three chapters sit in the middle of the book rather than in an appendix.

Chapter 55 . Replication, Preregistration, and the Falsification Suite

The Law

Paper E, §5.2, the minimal estimation model, quoted exactly:

```
MINIMAL ESTIMATION MODEL

    Y_{t+1} = a * Y_t  +  b * F(U_t)  -  Chat_t  +  eps_t

CONSTRAINTS

    0 <= a <= 1        (retention cannot exceed perfect retention,
                        and cannot be negative)
    b >= 0             (conversion efficiency cannot be negative)

IDENTIFICATION

    a = 1 - delta      ->   delta_hat = 1 - a_hat
    b = eta            ->   eta_hat   = b_hat

WHERE

    Y_t     an observed proxy for stored order Q_t
    U_t     the observed non-parallelism, in [0,1]
    F(U)    = ln(1 + (N-1)U),  the surplus law, Paper A Theorem 2.1
    Chat_t  an independently measured cost series, NOT a free parameter
    eps_t   residual
```

Paper E, §5.3, the falsification template, quoted exactly:

```
PROCEDURE

    1.  PREREGISTER  the operationalisations, the window, the split,
        the baselines, and the falsification rules.
    2.  FIT          on the training window only.
    3.  FREEZE       every choice. No refitting, no re-specification,
                     no re-selection of proxy, window, or period length.
    4.  REPORT       held-out prediction against four baselines:
                     (i)   persistence
                     (ii)  direct-input (raw U, no log)
                     (iii) matched autoregressive, no interaction term
                     (iv)  SHUFFLED-INPUT CONTROL

FALSIFIED IF

    RULE 1.  the shuffled-input control predicts held-out Y as well as
             the ordered series;
    RULE 2.  the matched autoregressive baseline is not beaten;
```

```
    RULE 3.  the fitted a lies outside [0,1] or the fitted b is negative.
```

This chapter applies that template to **subjective** quantities: to Y read off a magnitude estimate rather than a headcount, to U computed over one person's own framings rather than over a committee, and to the Moment `M_t` standing in the input slot in place of `F(U_t)` by the normalisation of CANON_INTERNAL §1 (`M` mapped into `[0, ln N]` serves as `S_t`).

Nothing in the template changes when the quantities become subjective. Everything in the threat model does. That asymmetry is the whole chapter.

55.0.1 What is proven and what is postulated

Proven. The identification map is exact algebra. Given the storage law `Q_{t+1} = (1−delta)Q_t + eta∗S_t − C_t` (Paper E §3) and given `Y_t = Q_t`, matching coefficients gives `a = 1−delta` and `b = eta` with no residual freedom. The constraint set follows from the definitions: `delta` is a retention loss in `[0,1]`, so `a = 1−delta` is in `[0,1]`; `eta` is a conversion efficiency and is nonnegative by construction. The bound `0 <= F(U) <= ln N` is Paper D, Corollary 2.2. The exact solution (Theorem 3.1), the one-step growth condition (Corollary 3.2), and the steady state `Q_inf = (eta∗F(u∗) − c∗)/delta` (Theorem 3.5) are theorems and are used here only as consequences of the fitted parameters, never as evidence for them.

Postulated, badged. That `Y_t` stands in a stable, monotone, and ideally affine relation to `Q_t`. That `delta`, `eta`, `N`, and the operationalisation of u are approximately constant over the estimation window. That the relation is linear in `Y_t` and in `F(U_t)` at all. That the subject's rating scale is ratio-scaled in the sense of Chapter 53 and stable in the sense of Chapter 54. **Each of these is a modelling assumption, each is falsifiable by the procedure in this chapter, and none of them is a theorem.** Chapter 52's FVP-6 gate exists precisely so that these four are stated in print before any data are collected.

Honesty clauses, repeated. The triangle inequality fails for geometric triples (Paper A §5); subadditivity holds on u only. As deployed, the surplus functional is a LOWER BOUND on the ensemble's information — we log the floor, not the ceiling. Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D does not identify surplus with Shannon entropy or mutual information. Paper E proves no thermodynamic negentropy law: `F(u) + N_can(u) <= ln N` (Corollary 2.2), and higher static surplus does not imply higher canonical concentration.

55.0.2 What changes when Y is subjective

Four things, and they are the reason Volume Two needs its own falsification chapter rather than a cross-reference to Chapter 25.

```
1. THE INSTRUMENT IS THE SUBJECT.
  In Volume One the rater and the rated were different objects. Internally they are
  the same organism. The person who supplies Y_t also has a view about whether the
   theory is true.

2. THE RATING CARRIES ITS OWN NOISE, AND THAT NOISE BIASES delta UPWARD.
  Measurement error on a LAGGED regressor attenuates a_hat. Since delta = 1 −
a,
   attenuation on a inflates delta. §55.9 computes the size.
```

```
3. THE INPUT AND THE OUTPUT ARE OFTEN RECORDED IN THE SAME SITTING.
  Anything that moves both at once — mood, fatigue, the weather, the fact of being
   measured — is a common cause the shuffle cannot see.

4. MEMORY RECONSTRUCTS.
  A retrospective rating of Y_{t+1} is partly a rating of what the subject believes
  happened at t. That is a lagged dependence with the right sign and the right timing,
   and it is not the dependence the theory predicts.
```

Threat 4 is the dangerous one, because **it produces exactly the temporal structure the shuffled control is designed to detect, and therefore passes the control.** §55.8 measures how often. The answer, in the configuration simulated there, is 52.7 per cent of the time.

55.0.3 The three rules, restated for the internal instrument

```
RULE 1 (SHUFFLE)    If the input series randomly permuted in time predicts held-
out
                  Y as well as the input series in its true order, there is no
                      temporal relationship. Report that.

RULE 2 (AUTOREGRESSION) If a model containing only Y_t and an intercept predicts
                   held-out Y as well, then F(U_t) contributed nothing and the
                      surplus term is decoration.

RULE 3 (CONSTRAINTS) If a_hat leaves [0,1] or b_hat goes negative, the fitted object
                 is not a storage law. It is a regression with the same letters.
```

A framework that cannot fail cannot be trusted, and this one names exactly how it fails. The three rules are not hedges. They are the price of admission.

The Engineering

55.1 The model, slot by slot

What goes in each slot when the scope is internal:

Slot	External (Volume One)	Internal (Volume Two)	Chapter
`Y_t`	institutional order proxy	character, skill, or internalised structure, via magnitude estimate	40, 53
`U_t`	non-parallelism between members	non-parallelism between one's own framings	39
`F(U_t)`	group surplus `ln(1+(N−1)U)`	internal surplus `F_int`	39

Slot	External (Volume One)	Internal (Volume Two)	Chapter
M_t	the situation as constructed	the moment as lived, (D/P)E	34
Chat_t	coordination overhead	cognitive load, effort, maintenance of self	5, 40
N	distinct domains held in a group	distinguishable categories in one mind	7, 32

Two admissible input series exist internally, and they must not be mixed inside one preregistered model:

```
INPUT VARIANT A   S_t = F(U_t) = ln(1 + (N−1) U_t)
                  U_t measured over one person's own framings (Chapter 39)

INPUT VARIANT B   S_t = M_t normalised into [0, ln N]   (CANON_INTERNAL §1)
                  M_t = (D_t / P_t) ∗ E_t                (Chapter 34)
```

Variant B inherits everything Chapter 53 says about ratio scales. **If D and P are not ratio-scaled with a true zero, D/P is not a number and Variant B is not a model.** The preregistration must state which variant is being tested, and it must state the evidence that the scale requirement is met, or demote the formula in print as FVP-2 requires.

Chat_t is not a free parameter. It is an independently measured series, subtracted before fitting. **Allowing C to be estimated inside the regression is the single most common way this model is made unfalsifiable**, because a free additive term with an intercept and a trend will absorb any misfit that has the shape of a level or a drift. Fit two coefficients. Measure the third.

55.2 The constraint set is the first falsification rule, not a convenience

It is tempting to fit by ordinary least squares and then note with satisfaction that a_hat came out inside [0,1]. That gets the logic backwards. The constraints are the theory. a outside [0,1] means the system amplifies its own stored order period on period without input, which is not a leaky integrator but an oscillator or a divergence. b < 0 means interaction destroys stored order, which contradicts eta >= 0 in the definition of the storage law. **Fit the model constrained; report the unconstrained fit alongside it; and if the unconstrained fit lands outside the box, say so in the abstract, because that is Rule 3 firing.**

There is a practical trap. A constrained fit cannot report a violation it is not allowed to express. If a_hat sits exactly at the boundary — a_hat = 1.000 or b_hat = 0.000 — the optimiser has hit the wall, and a wall-hit is evidence, not a result. In the null simulation of §55.7 the constrained fit returned a_hat = 1.0000 exactly: retention pinned at unity, no decay at all, which is the fit telling you it wanted to leave the box.

```
BOUNDARY REPORTING RULE

    Report a_hat and b_hat from the CONSTRAINED fit.
    Report a_ols and b_ols from the UNCONSTRAINED fit.
    Flag in bold any of:
```

```
        a_hat = 0 or a_hat = 1 exactly
        b_hat = 0 exactly
        a_ols outside [0,1]
        b_ols < 0
    A flagged boundary is a Rule 3 event and must appear in the abstract.
```

55.3 The preregistration document

Preregistration is a device that removes a specific, enumerable list of choices from the researcher's hands after the data arrive. Each line below is a degree of freedom that, left open, can manufacture a result.

```
PREREGISTRATION — REQUIRED FIELDS

A. SUBJECT AND SCOPE
   A1  scope: internal (single person) / dyadic / group
   A2  number of subjects, and whether analysis is per-subject or pooled
   A3  inclusion and exclusion rules, stated before recruitment

B. THE PERIOD
   B1  period length (a day, a week, a fortnight) and the justification
   B2  the alignment rule: what counts as period t for Y, for U, for C
   B3  the handling of missed periods (interpolate? drop? carry forward?)

C. OPERATIONALISATIONS
   C1  Y: instrument, items, scoring, and the argument that Y is not a
       repackaging of U (Paper E §6.3 — the most common measurement error)
   C2  U: the framing set, the inner-product definition, the computation of u
   C3  N: how the category count was audited, not guessed (Chapter 7)
   C4  C: the cost instrument and its units, in the same units as eta*S
   C5  input variant A or B, and for B the ratio-scale evidence (Chapter 53)

D. THE SPLIT
   D1  T_train and T_holdout, fixed in advance as counts, not as fractions
       of a total that is not yet known
   D2  the split is CHRONOLOGICAL. No random splits on time series.
   D3  no peeking rule: who holds the held-out data and when it is opened

E. THE BASELINES
   E1  persistence
   E2  direct-input (raw U, no log), with the required u-range (§55.5.2)
   E3  matched autoregressive with intercept
   E4  shuffled-input control: family, B, and the calibration (§55.7)

F. THE FALSIFICATION RULES
   F1  Rule 1 threshold, as a calibrated p-value, not a nominal one
   F2  Rule 2 threshold, as a skill score with its sign convention
   F3  Rule 3 boundary reporting rule
```

```
G. THE STOPPING RULE
   G1  when data collection ends, stated as a date or a count, never as
       "when the result stabilises"

H. THE PUBLICATION COMMITMENT
   H1  the commitment to report the result whichever way it falls
   H2  the deposit location for the frozen model and the analysis code
```

A stopping rule that reads "until the pattern is clear" is not a stopping rule; it is optional stopping, and optional stopping will eventually produce any pattern you like. Fix the count. Fix the date. Publish whichever way it falls.

55.4 The freeze

"Freeze" means: after the training fit, the following are constants and may not be touched again.

```
FROZEN AT THE END OF TRAINING

    a_hat, b_hat                      the two coefficients
    the operationalisation of Y       items, scoring, reverse-keying
    the operationalisation of U       framing set, inner product
    N                                 the category count
    the cost instrument               and its units
    the period length                 and the alignment rule
    the missing-data rule
    the input variant                 A or B
    the transformation of Y           log, none, winsorisation limits
    the baseline definitions          all four
    the surrogate family and B        for the shuffle
    the skill-score definition
    the falsification thresholds
```

The list is long because researcher degrees of freedom are numerous and each is individually defensible. **Every item on it has a reasonable-sounding argument for changing it after the data are seen, and that is exactly why it must be frozen before.**

One permitted exception, declared in advance: a data-quality exclusion rule that does not reference Y, U, or the fit. "Periods in which the subject was hospitalised are excluded" is legitimate if written in field A3 before collection. "Periods in which the residual exceeded two standard deviations are excluded" is not legitimate under any circumstances.

55.5 The four baselines

A baseline is a rival explanation with a name. Each of the four names a specific way the result could be uninteresting.

55.5.1 Persistence

```
    Y_{t+1} = Y_t
```

No parameters. Rival explanation: *the thing you measured is simply slow-moving.* Persistence is the hardest baseline to beat for any highly autocorrelated quantity, which is most psychological constructs at daily or weekly cadence. It also tells you whether your period length is sensible: **if persistence is unbeatable, your period is shorter than the dynamics you are trying to observe, and the fix is a longer period, not a better model.** In the worked example it gives a held-out RMSE of 0.1207 against the model's 0.0777, a skill score of 0.5860.

55.5.2 Direct-input

```
Y_{t+1} = a * Y_t + b * U_t - Chat_t
```

Raw u, no logarithm. Rival explanation: *the log in Paper A's F is doing no work; any increasing function of u would fit as well.*

This baseline tests Paper A Theorem 2.1's functional form, and it is the weakest of the four for a reason worth computing. Over u in `[0,1]` with `N = 7`, the correlation between u and `F(u)` is 0.9724. The best straight-line approximation is `F(u) ~= 1.7280 u + 0.4063` with a maximum absolute error of 0.4021 and an RMS error of 0.1197 — visible over the full range. But over any narrow band the two are nearly identical:

```
u range        corr(u, F(u))
[0.0, 0.2]        0.99496
[0.2, 0.5]        0.99704
[0.5, 1.0]        0.99741
[0.3, 0.6]        0.99795
```

The direct-input baseline has discriminating power only in proportion to the range of u your data actually span, and most field data span very little of it. Simulated at `T = 100` with 70 training periods, the fraction of trials in which the log form beats the linear form on held-out data rises with the observed span of u:

```
mean span of u        log form wins
     0.323                84.0 %
     0.580                91.3 %
     0.862                97.7 %
     0.987               100.0 %
     1.000               100.0 %
```

and the ratio of linear-model to log-model held-out RMSE rises from 1.063 at span 0.338 to 1.606 at span 0.987.

The consequence for preregistration is concrete. Field E2 must state a **minimum observed span of u** below which the direct-input comparison is declared uninformative and reported as such. A defensible default from the table above is a span of 0.60. Below that, say in print: *the data do not distinguish the logarithmic from the linear input form.* That is an honest partial result and it is much better than a claim the data cannot support.

55.5.3 Matched autoregressive

```
Y_{t+1} = a * Y_t + c - Chat_t
```

An intercept in place of the surplus term. Rival explanation: *the system decays toward a constant and the interaction variable is irrelevant.* This baseline isolates the theory's actual content, because the storage law without `eta*S_t` is exponential decay to a level. It is also the baseline most likely to win when the input barely varies: **if `F(U_t)` is nearly constant, its coefficient is unidentifiable and an intercept does the same job at lower cost.** In the null simulation of §55.7 Rule 2 fires — model RMSE 0.0978 against matched-AR 0.0681, skill -1.0669.

But Rule 2 alone is nowhere near sufficient. In 500 simulated null trials — data generated with no dependence on S at all — **the model beat the matched autoregressive baseline on held-out data 35.8 per cent of the time.** A single baseline comparison, run once, on data where both series are smooth, is a coin weighted only slightly against you.

55.5.4 The shuffled-input control

```
Take the input series S = (S_0, S_1, ..., S_T).
Permute it in time to get S' = (S_{pi(0)}, S_{pi(1)}, ..., S_{pi(T)}).
Refit on training with S'. Score on held-out with a permuted held-out S'.
Repeat B times. Compare the ordered model's held-out error to the
distribution of the permuted models' held-out errors.
```

Rival explanation: *the input and the output have the right marginal distributions and no temporal relationship at all.*

The shuffle preserves the marginal distribution of the input exactly — every value, its mean, its variance, its skew, its range — and destroys only the pairing between input and output in time. That is the precise null it instantiates, and it is the reason it is the sharpest of the four.

Everything else that could be responsible for a fit survives the shuffle: the scale of S, the number of large values, the fact that S is bounded by `ln N`, the fact that Y trends, the fact that both are positive. The only thing the shuffle removes is *which* S went with *which* Y. **If the fit survives that removal, the relationship you found was a property of the two marginal distributions and not a relationship between two series.**

55.6 Why the shuffled control is the one most often omitted

The other three baselines are, in practice, strawmen: the researcher expects to beat persistence (the construct moves), direct-input (the log is in the theory), and matched autoregression (the input term is in the theory). Running them feels like diligence and usually confirms what was already believed.

The shuffle is different in kind. **It is the only one of the four that can kill a result the researcher already believes, using data the researcher already has, at no additional cost, in seconds of computation.** That combination — decisive, cheap, already possible — is what produces omission, because there is no resource-based excuse to hide behind. Its omission is a choice about what one wishes to know.

Three reasons it gets left out, in the order of how often they are offered:

1. `"It is not standard in this literature."`
   `Irrelevant. A permutation test needs a random number generator, not a literature's permission.`

2. "The input is autocorrelated so the shuffle is not a fair comparison." A real technical objection, answered by choosing the right surrogate family (§55.7), not by dropping the control.

3. "We already showed the effect is significant." Significance against a parametric null assumes the null. The shuffle CONSTRUCTS the null from the data you have. When the two disagree, believe the shuffle: it assumed less.

There is a fourth reason, unstated: the shuffle is the only baseline whose result the researcher cannot predict in advance. **A test whose outcome you can predict is not a test.**

55.7 The surrogate hierarchy, and why you must calibrate the control itself

The objection in reason 2 is correct, and taking it seriously produces the most technically important section of this chapter.

An independent-and-identically-distributed permutation destroys two things at once: the pairing between S and Y (which you want destroyed) and the autocorrelation within S (which you did not). If both S and Y are smooth, a smooth ordered S can track a smooth Y by coincidence, while the whitened shuffled S cannot track anything. **The plain shuffle can therefore reject the null when the null is true, because it compares a smooth regressor against a white one rather than against another smooth one.** This is the spurious-regression problem wearing a permutation test's clothes.

Three surrogate families, in increasing fidelity to the input's own structure:

```
IID PERMUTATION      S' = random permutation of S
                preserves: the multiset of values, hence every marginal moment
                     destroys:  pairing AND autocorrelation

BLOCK PERMUTATION   cut S into contiguous blocks of length L; permute the blocks
                     preserves: marginals, and autocorrelation at lags < L
                     destroys:  pairing, and autocorrelation at lags >= L
                   choose L >= 3 / (-ln rho1) so a block spans the input's own
                     correlation time

CIRCULAR SHIFT       S' = S rotated by a random offset k
                   preserves: marginals AND the full autocorrelation function
                     destroys:  pairing only
                  the tightest surrogate; the number of distinct surrogates is
                     limited to the series length, which caps resolution
```

Now the measurement that decides the matter. Five hundred null trials were simulated: Y generated by a storage law driven by a smooth AR(1) input that is *not* S, with S an independent smooth AR(1) series (autocorrelation parameter 0.85 for both), `T_train = 40`, `T_holdout = 20`, 200 surrogates per trial. The true b is zero. A well-calibrated test at nominal alpha = 0.05 should reject 5 per cent of the time.

```
NULL (b_true = 0), 500 trials, nominal alpha = 0.05
```

```
   surrogate family      measured rejection rate
   iid permutation              0.108
   circular shift               0.120
   block permutation (L=5)      0.060

   matched-AR comparison alone: passes 0.358 of the time
   mean b_hat under the null:   0.0922  (sd 0.0222)

TRUE MODEL (b_true = 0.40), 500 trials

   surrogate family      measured rejection rate (= power)
   iid permutation              1.000
   circular shift               1.000
   block permutation (L=5)      0.998

   beats matched AR on held-out: 1.000
   mean b_hat:                   0.4025  (sd 0.0322)
```

Read those two tables together and three conclusions follow, none of which is optional.

First: the shuffled-input control, run as it is usually run, is anti-conservative. At a nominal 5 per cent it rejected between 6.0 and 12.0 per cent of true nulls in this design. A control that is wrong by a factor of two in the direction of false discovery is still enormously better than no control — compare the 35.8 per cent false pass rate of the matched-AR comparison alone — but it is not exact, and reporting its nominal p-value as though it were exact is a misstatement.

Second: the block permutation was the best calibrated of the three here (0.060 against a nominal 0.050), and it lost almost nothing in power (0.998 against 1.000). That is a favourable trade and it should be the default. Note that the circular shift, which preserves the most structure, was *not* the best calibrated in this design, because the held-out window is short and rotating a 21-point segment leaves few genuinely distinct surrogates. Intuition about which surrogate is "tightest" does not survive contact with a short held-out window.

Third, and this is the specification item: calibrate the control on your own design before you trust its p-value. The false-positive rate of a surrogate test depends on the autocorrelation of both series, the length of the training and held-out windows, the number of surrogates, and the constraint set. None of those is universal. The calibration is a simulation you run once, in advance, using the preregistered design and a generating process with b set to zero, and it takes minutes. The output is a *calibrated* critical value that replaces the nominal one in field F1.

Finally, resolution. A permutation p-value cannot be smaller than `1/(B+1)`:

```
   B        smallest attainable p     se of p at p = 0.05
   19            0.05000                   0.05000
   99            0.01000                   0.02190
  199            0.00500                   0.01545
  999            0.00100                   0.00690
 1999            0.00050                   0.00487
```

`B = 199` is the practical floor for a claim at the 1 per cent level; `B = 999` if the claim is stronger. Report B. A shuffle test whose B is not reported is not reproducible.

55.8 What the shuffled control does not catch

Here is the honest limit, and it is severe.

The shuffle detects the absence of a temporal relationship. It cannot distinguish a real temporal relationship from a fabricated one that has the same timing. Any process that writes the input's timing into the output — by any route other than the theorised one — will produce a series that passes the shuffle handsomely.

The mechanism that does this most reliably is threat 4 of §55.0.2: a subject who knows the hypothesis and whose rating of `Y_{t+1}` is nudged by what they recall of period `t`. Simulated directly: data generated with `b_true = 0` — no real dependence whatsoever — and the *reported* `Y_{t+1}` inflated by `0.45 * (S_t - mean S)`.

```
DEMAND-CHARACTERISTIC SIMULATION, true b = 0, T_train = 40, T_holdout = 20

  example trial   fitted a = 0.9265   fitted b = 0.1856   (both inside the box)
                   held-out RMSE  model 0.0991
                                  matched AR 0.1112
                                  circular-shift surrogate mean 0.1168
                   shuffle p = 0.0000

  over 300 trials, the full four-baseline suite PASSED a purely fabricated
  relationship 52.7 % of the time
```

More than half the time, a relationship that does not exist cleared all four baselines and all three rules, because the rater's own bias supplied the timing. No amount of extra shuffling fixes this. The shuffle is testing the wrong thing: the fabrication is temporal, so a temporal test endorses it.

The fix is not statistical. It is a design change, and there are four, in decreasing order of strength:

```
D1  BLIND THE RATER TO THE HYPOTHESIS.
    The subject supplies Y without being told what U is, that U is being
    measured, or what relationship is predicted. Debrief afterwards.

D2  MEASURE Y AND U THROUGH DIFFERENT CHANNELS.
    Y from a behavioural or archival trace (a completed artefact, an
    attendance record, a produced object); U from the rating instrument,
    or the other way round. A bias in one channel cannot supply timing
    to the other.

D3  PROSPECTIVE, NOT RETROSPECTIVE, SAMPLING.
    Rate the period in the period, not at the end of the week. Memory
    reconstruction is proportional to the recall interval; shrink the
    interval and you shrink the artefact.

D4  YOKED CONTROL SUBJECT.
    A second subject is shown the first subject's U series, on the same
    schedule, but their own Y is their own. If the yoked subject shows
    the same fit, the fit is being produced by the SCHEDULE, not by the
```

```
subject's own dynamics.
```

D1 and D2 together are close to sufficient. **State in the preregistration which of D1-D4 are in force. A study running none of them cannot claim a temporal result about subjective quantities regardless of how many surrogates it ran.**

A broader lesson belongs in the manual's spine. Chapter 25 taught that the shuffle is the decisive baseline, and at group scope with archival proxies it very nearly is. Turned inward, **the shuffle remains the most decisive of the four and stops being sufficient on its own.** The second half is what Volume Two adds.

55.9 The instrument's own noise, and which direction it biases

A subjective Y is a rating, and a rating has error. That error does not merely add scatter; it biases both parameters, in known directions, by computable amounts.

The mechanism is regression dilution applied to a lagged regressor. Error in `Y_t` attenuates `a_hat` toward zero, and since `delta_hat = 1 − a_hat`, **attenuation of a inflates `delta`: measurement noise makes a system look leakier than it is.** Because the lagged term then explains less, the input term compensates, so `b_hat` — hence `eta_hat` — is inflated too.

Simulated at `T = 80`, true `a = 0.85`, true `b = 0.40`, process noise 0.06, 400 replications per row:

rating sd	mean a_hat	mean b_hat	implied delta_hat
0.00	0.8486	0.4022	0.1514
0.05	0.8419	0.4156	0.1581
0.10	0.8219	0.4538	0.1781
0.20	0.7571	0.5783	0.2429
0.40	0.5683	0.9425	0.4317

Read the last row. With a rating standard deviation of 0.40 against a Y whose own standard deviation is around 0.20, the estimated `delta` is 0.432 against a true 0.150 — nearly triple — and the estimated `eta` is 0.943 against a true 0.400, more than double. **A sloppy instrument does not produce a noisy answer. It produces a confidently wrong one, and it errs in the direction of making the system look both leakier and more efficient than it is.**

This is why Chapter 56's Stage 2 comes before Stage 3 and cannot be skipped: the scale must be built and its error characterised before any dynamic parameter estimated on it means anything. It is also why the preregistration must report the instrument's test-retest reliability alongside the parameter estimates. **A parameter estimate without an accompanying reliability figure for the instrument that produced it is uninterpretable, because the reader cannot tell how much of the reported `delta` is the world and how much is the ruler.**

A partial correction exists — an attenuation factor estimated from a repeated-measures reliability study — but it depends on that estimate being itself accurate and does nothing for `b_hat` without further assumptions. **Prefer a better instrument to a correction factor.**

55.10 Held-out reporting

Report one table. It has five rows and it is the entire empirical claim.

```
HELD-OUT REPORT (the only table that matters)

  predictor                        RMSE      skill vs model      note
  ------------------------------------------------------------------------
  frozen model                    ____         --
  persistence                     ____        ____
  direct-input (raw u)            ____        ____             span of u = ___
  matched autoregressive          ____        ____
  shuffled input (family, B)     ____        ____           calibrated p = ___
```

Definitions, fixed in advance:

```
  RMSE         root mean squared ONE-STEP-AHEAD error on the held-out window,
               using OBSERVED Y_t on the right-hand side, never simulated Y_t

  skill score   1 - (RMSE_model / RMSE_baseline)^2
                positive means the model beats the baseline;
                zero means they tie; negative means the baseline wins
```

Two prohibitions that matter more than they look.

Prohibition one: one-step-ahead scoring, using observed `Y_t`. Free-running simulation from `Y_0` is a weaker test that mostly measures whether `Q_inf` was estimated well; as a primary score it inflates apparent performance for any model with a stable fixed point. Secondary diagnostic only.

Prohibition two: no refitting. If the held-out window is opened and the model performs poorly, the result is that the model performs poorly. A revised model fitted after seeing held-out data is a new hypothesis, requires new held-out data, and must be labelled exploratory in every subsequent sentence about it. **The held-out window can be spent exactly once.**

Report the in-sample fit too, and report it in a way that makes its irrelevance visible. In the null simulation of §55.7 the training-window R^2 was 0.6477 — a figure that would be reported as a strong result in a great many papers — while the same frozen model was beaten on held-out data by persistence, by matched autoregression, and by 100 per cent of the circular-shift surrogates. **In-sample R^2 of 0.65 and a held-out failure against every baseline is not a contradiction. It is the normal appearance of a false result.**

55.11 Cross-individual parameter recovery — FVP-7, the hard gate

Everything to this point tests whether a relationship exists in one series. FVP-7 asks the question that decides whether the framework is science: **do `delta` and `eta` recover as properties of a person, measured the same way, by different investigators?**

CANON_INTERNAL §6 identifies this as the gate most often skipped, and skipping it is what separates a formalism that describes from one that measures. Chapter 54's specific objectivity is the standard aimed at — a person's measure must not depend on which items were used — and the analogous demand here is that a person's `eta` must not depend on which lab estimated it. Three levels, cleared in order.

Level 1 — within-person reliability Split one subject's series into two halves and estimate separately in each. If the halves disagree, the parameter is a property of the window, not of the person. Simulated with rating noise 0.10 and true between-person spread 0.12 in b, 40 subjects per row:

```
half length      split-half r of b_hat
     60                  0.841
    120                  0.846
    240                  0.962
```

Sixty periods per half is enough for a usable within-person reliability; 240 is enough for a good one. Below 60, do not attempt Level 2.

Level 2 — between-person separability Estimate for many subjects and ask whether the spread exceeds what estimation error alone produces. This step is almost universally botched, because the spread is always nonzero and always tempting to interpret. Simulate the noise floor directly: give every subject the *same* true `b = 0.40` and see how much apparent individual difference the estimator manufactures out of nothing.

```
THE NOISE FLOOR — 60 subjects who genuinely share one eta

T per subject      sd of b_hat across subjects
     60                   0.0478
    120                   0.0355
    240                   0.0255
    360                   0.0166
```

At 120 periods per subject, an estimator applied to sixty identical people reports a between-person standard deviation of 0.036 in `eta`. Any story about individual differences built on a spread smaller than that is a story about the estimator. A separate 30-subject run with the between-person standard deviation set to exactly zero produced an apparent spread of 0.042 — the same conclusion from a different seed.

This is why the recovery correlation collapses when true differences are small. With 30 subjects at 120 periods each:

```
true between-person sd of b        corr(true b, estimated b)
           0.12                            0.956
           0.04                            0.503
```

The second row is not a failure of the estimator. **The second row is what it looks like when the quantity you are trying to measure varies less across people than your instrument's error, and no sample size fixes it — only a longer series per person or a quieter instrument does.**

The full variance decomposition for 30 subjects at 120 periods, true between-person sd 0.12:

```
var(true b)          0.011586
var(estimated b)     0.012707
var(error)           0.001310
reliability          var_true / (var_true + var_err) = 0.898
```

Reliability 0.898 caps the correlation between `eta` and any external variable at `sqrt(0.898)` = 0.948. That ceiling is the useful output: **before correlating `eta` with anything — a personality**

measure, an outcome, a treatment response — compute the reliability and state the ceiling, because a correlation cannot exceed it and a reported correlation near it is suspicious.

Also report the bias. Across the recovery runs, `b_hat` was biased upward by between 0.038 and 0.078, which is the attenuation mechanism of §55.9 operating at rating noise 0.10. **The bias is in the mean, not in the ranking, which is why recovery correlations can look excellent while every individual estimate is too high.** Report both.

Level 3 — between-lab invariance The hard gate proper. Two labs measure the same construct with operationalisations they each consider faithful. Lab 2's procedure for computing u runs 15 per cent hot — a slightly wider framing set, a slightly different inner product, an entirely defensible choice.

```
lab 1 mean b_hat        0.4679
lab 2 mean b_hat        0.4068
observed ratio          0.8696
predicted ratio 1/1.15  0.8696
```

Exact, because the model is linear in the input: scaling S by g scales `b_hat` by `1/g`. That gives a useful diagnostic:

```
BETWEEN-LAB DIAGNOSTIC

If lab 2's eta differs from lab 1's by a CONSTANT FACTOR across all subjects,
the discrepancy is a CALIBRATION difference in the operationalisation of u,
and it is repairable by publishing the scaling constant.

If lab 2's eta differs from lab 1's by a factor that VARIES ACROSS SUBJECTS,
or if the RANK ORDER of subjects differs between labs, the discrepancy is a
THEORY FAILURE. The parameter is not a property of the person.
```

A framework passes FVP-7 when between-lab disagreement is a single number and fails when it is a pattern. That is a sharp, checkable, preregisterable criterion, and it is the criterion Chapter 56 Stage 4 is built to apply.

Predicted ratios for other calibration errors, all exact:

```
inflation g     predicted b_hat ratio 1/g
   1.05                 0.9524
   1.10                 0.9091
   1.15                 0.8696
   1.25                 0.8000
   1.50                 0.6667
```

55.12 Goodhart and Campbell

Goodhart's Law: when a measure becomes a target it ceases to be a good measure. Campbell's Law: the more a quantitative indicator is used for social decision-making, the more it is subject to corruption pressures and the more it distorts what it was meant to monitor.

Both attach to Y, the order proxy, and the attachment is computable. Simulate a system running normally for 30 periods, at which point Y becomes a target and the *reported* value inflates by 0.35 while the underlying Q continues to obey the same law.

```
GOODHART SIMULATION — proxy becomes a target at t = 30

at t = 30    true Q = 3.0359      reported Y = 3.3650
at t = 60    true Q = 2.8447      reported Y = 3.1401   inflation = 0.2954

fitted on periods  0–30 (pre-target):   a = 0.8992  b = 0.3105
                                        -> delta = 0.1008  eta = 0.3105
fitted on periods 30–60 (post-target):  a = 0.8817  b = 0.3551
                                        -> delta = 0.1183  eta = 0.3551
```

The estimated conversion efficiency rose by 14 per cent — from 0.3105 to 0.3551 — **at exactly the moment the underlying stored order began to fall, from 3.036 to 2.845.** Nothing in the fit complains. The constraints hold. The parameters stay in the box. The model reports an improvement while the thing improves in the opposite direction.

Goodhart's Law is not a caveat about this framework. It is a failure mode with a signature, and the signature is a rising `eta_hat` accompanied by a falling independent indicator. That gives the detection procedure:

```
GOODHART DETECTION

G1  Maintain at least one order indicator that is NEVER reported upward,
    never used in any decision, and known to no one who can influence it.
    This is the shadow proxy. It exists solely to be compared.

G2  Test for a structural break in (a_hat, b_hat) at every point where the
    measure acquired a consequence: a review cycle, a published dashboard,
    a funding decision, a self-imposed goal.

G3  Compare the trajectory of the reported proxy against the shadow proxy.
    A widening gap dated to a decision is Campbell's Law, measured.

G4  Prefer proxies that are EXPENSIVE TO FAKE relative to the benefit of
    faking them. Archival traces and produced artefacts beat self-report
    on this criterion by a wide margin.
```

Internally — one person, measuring themselves, against their own goals — the problem is at its worst, because **the subject, the instrument, the analyst, and the beneficiary of a good result are one organism with one set of incentives.** Volume One could at least separate the rater from the rated. Here the separation must be manufactured: by D1-D4 of §55.8, by the shadow proxy of G1, and by the honest expectation that self-measurement drifts toward flattery at a rate that is itself worth measuring.

CANON §5 refusal 4 applies here with unusual force: do not build a measurement regime that removes a person's ability to decline being measured. **A proxy that a person cannot refuse to feed is a proxy they will learn to feed falsely, and the falsity will be indistinguishable from success.**

55.13 The publication contract

P1 Publish the result whichever way it falls. A preregistration that is abandoned when the result disappoints has produced nothing except a private lesson.

P2 Deposit the frozen model — the two coefficients, the code, the operationalisations, the surrogate calibration — at the moment of the freeze, before the held-out window is opened.

P3 Report B, the surrogate family, and the CALIBRATED critical value.

P4 Report the instrument's test-retest reliability in the same table as the parameter estimates (§55.9).

P5 Report the observed span of u, so the reader can judge the direct-input comparison (§55.5.2).

P6 Report boundary hits in the abstract (§55.2).

P7 If any of the three rules fired, say so in the first paragraph. Not the limitations section. The first paragraph.

P8 If the result is null, the title should say so.

P8 is not moralising. **A field in which nulls are titled ambiguously is a field whose literature cannot be read at face value, and Volume Two's premise is a claim about literatures, not about individual papers.** The Open Science Collaboration's 2015 replication project (*Science*) found 97 of 100 originals significant, about 36 per cent of replications significant, effect sizes roughly halved. **Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does** — and so does a convention under which a null is legible from the title.

The Specification

A study that cannot answer yes to all of these has not run the falsification suite; it has run a regression.

S55.1 The preregistration contains every field A1–H2 of §55.3, filed before the first observation, with a timestamp that a third party can verify.

S55.2 The input variant (A: F(U_t), or B: normalised M_t) is declared, and if B, the Chapter 53 ratio-scale evidence for D and P is filed with it.

S55.3 Chat_t is measured independently and subtracted before fitting. It is never estimated as a free parameter. Its units match those of eta∗S.

S55.4 The split is chronological, with T_train and T_holdout fixed as counts in advance.

S55.5 The model is fitted with the constraints 0 <= a <= 1 and b >= 0, and the unconstrained fit is reported alongside it.

S55.6 Boundary hits (a_hat at 0 or 1, b_hat at 0, or any unconstrained estimate outside the box) are flagged in the abstract as Rule 3 events.

S55.7 All four baselines are computed on the held-out window: persistence, direct-input, matched autoregressive, shuffled input.

S55.8 The surrogate family is declared in advance. Block permutation with L >= 3 / (-ln rho1) is the default; departures are justified in writing.

S55.9 B >= 199 surrogates for a claim at the 1 per cent level; B >= 999 for anything stronger. B is reported.

S55.10 The surrogate test is CALIBRATED by simulation on the preregistered design with b set to zero, and the calibrated critical value replaces the nominal alpha in field F1. Nominal p-values are not reported alone.

S55.11 The observed span of u on the training window is reported. If it is below the preregistered minimum (default 0.60), the direct-input comparison is declared uninformative in print.

S55.12 At least two of the four design controls D1–D4 (§55.8) are in force, and which two is stated. Blinding the rater to the hypothesis (D1) and separating the channels for Y and U (D2) are the preferred pair.

S55.13 The instrument's test-retest reliability is estimated in a separate study, reported in the same table as the parameter estimates, and used to state the ceiling on any correlation involving those estimates.

S55.14 Scoring is one-step-ahead using observed Y_t. Free-running simulation may be reported as a secondary diagnostic and never as the primary score.

S55.15 Skill scores use 1 - (RMSE_model/RMSE_base)^2 with the sign convention stated, and are reported for all four baselines.

S55.16 The held-out window is opened exactly once. Any model revision after that opening is labelled exploratory in every subsequent sentence.

S55.17 Within-person split-half reliability of a_hat and b_hat is reported, with at least 60 periods per half.

S55.18 Before any claim about individual differences, the estimation noise floor is simulated at the study's own T and n, and the observed between-person

spread is compared against it explicitly.

S55.19 The reliability var_true/(var_true + var_err) is computed and the implied ceiling on external correlations is stated.

S55.20 Between-lab comparison, where available, is reported as either a single scaling constant (calibration difference, repairable) or a pattern (theory failure, not repairable). The distinction is made in print.

S55.21 A shadow proxy exists (Goodhart G1): at least one order indicator never reported upward and never used in any decision.

S55.22 A structural-break test on (a_hat, b_hat) is run at every date on which the measure acquired a consequence.

S55.23 The frozen model and analysis code are deposited at the moment of the freeze, before the held-out window is opened.

S55.24 The result is published whichever way it falls, and if null, the title says so.

S55.25 No participant is required to be measured. Refusal is available, costless, and recorded as missing under the preregistered missing-data rule (CANON §5, refusal 4).

The Failure Modes

Each fault is named, given its diagnostic signature, and given its fix.

F55.1 The Free Cost Term

Signature. Chat_t appears as an estimated coefficient, or an intercept plus trend, rather than a measured series. In-sample fit excellent, held-out fit mediocre, residuals structureless because the free term absorbed everything. Measuring cost is tedious; estimating it is easy.

Fix. Measure C (Chapter 5, Appendix C.5). Subtract it. Fit two coefficients only. If C genuinely cannot be measured, say so and report that the model was fitted without a cost term, which is a different and weaker model.

F55.2 Proxy Contamination

Signature. Y rises whenever interaction is high, for reasons internal to the instrument. Paper E §6.3 names this as the most common measurement error in the entire framework: choosing an order proxy that is a repackaging of the interaction variable.

Fix. Field C1's argument, written before collection. Ask: could Y rise in a period with high U and no durable change whatsoever? If yes, Y is measuring U. Use D2 of §55.8 — different channels

for the two series.

F55.3 The Uncalibrated Shuffle

Signature. A reported permutation p-value of 0.02 with no statement of the surrogate family, no `B`, and no calibration run. The permutation test is usually presented as exact; it is exact only for the exchangeability structure it assumes.

Fix. §55.7. Simulate the design at `b = 0`, measure the actual rejection rate, and use the calibrated critical value. In the design simulated there, nominal 0.05 corresponded to actual rates between 0.060 and 0.120 depending on the surrogate family.

F55.4 The Fabricated Timing

Signature. Every baseline beaten decisively, shuffle p-value 0.000, parameters comfortably inside the box — and the subject knew the hypothesis, rated both series, and rated them retrospectively. It is the cheapest study to run.

Fix. D1-D4 of §55.8, at least two in force. Measured rate at which this fault passes the full suite when uncontrolled: **52.7 per cent.** There is no statistical remedy; only a design remedy.

F55.5 Attenuation Mistaken for Decay

Signature. `delta_hat` comes out large — the system looks leaky — and the instrument's reliability was never reported. Error in a lagged regressor attenuates its coefficient, and `delta = 1 − a`.

Fix. §55.9's table. Estimate instrument reliability separately. If rating error is comparable to the signal's own variability, the reported `delta` may be inflated by a factor approaching three (0.150 true against 0.432 estimated at rating sd 0.40).

F55.6 The Uninformative Direct-Input Comparison

Signature. "The logarithmic form outperformed the linear form" reported from data spanning `u` in `[0.30, 0.45]`. Over narrow bands `corr(u, F(u))` exceeds 0.995 and the comparison is near-vacuous.

Fix. S55.11. Report the span. Below the preregistered minimum, say the data do not distinguish the forms.

F55.7 Optional Stopping

Signature. The stopping rule reads "until the pattern is clear", or the reported `T` is not the preregistered `T`, or a supplementary wave was added after a first look. Stopping early when things look good feels efficient.

Fix. Field G1, as a date or a count. Any deviation is reported as a deviation, in the abstract.

F55.8 The Second Held-Out Window

Signature. The held-out result was poor; the model was revised; a new held-out window appeared from somewhere; the revised model is reported as confirmatory. The revision is usually a genuine improvement, which is what makes the fault feel like progress.

Fix. S55.16. The revised model is exploratory, permanently, until tested on data collected after its specification was frozen.

F55.9 Individual Differences That Are Estimation Noise

Signature. A table of per-subject `eta` values with a spread of 0.03 to 0.04, at 120 periods per subject, interpreted as a personality finding. The spread is real in the sense that the numbers differ, not in the sense that the people do.

Fix. S55.18. The measured noise floor at 120 periods is a between-subject sd of 0.036; at 60 periods, 0.048. A spread that does not clear the floor is the estimator talking.

F55.10 Goodhart Drift

Signature. `eta_hat` rising over successive fits while an independent indicator falls, the rise dated to a review cycle, a dashboard, or a personal goal. The measure acquired a consequence.

Fix. G1-G4 of §55.12. Simulated magnitude: `eta_hat` up 14 per cent (0.3105 -> 0.3551) while true `Q` fell from 3.036 to 2.845.

F55.11 The Unfalsifiable Rescue

Signature. After a failure, the model is retained with a new auxiliary assumption: `delta` was time-varying, `N` changed mid-study, the period length was wrong, the subject was unusual. Each rescue is individually plausible, which is precisely the problem.

Fix. Every auxiliary assumption becomes a preregistered prediction in the next study or it is withdrawn. A rescue that generates no new testable prediction converts a research programme into a system of excuses, and Chapter 56 §56.10 gives the tripwires.

F55.12 The Missing Null

Signature. A research group with five preregistrations and two publications. Nulls are harder to place and less rewarding to write.

Fix. P1 and P8. Publish. Title it.

Worked Example

One subject. Sixty-one periods of daily observation. `N = 7` distinguishable framings audited by the Chapter 7 procedure, so `ln N = 1.945910`. Input variant A: `S_t = F(U_t) = ln(1 + 6 U_t)`. `Y_t` is a magnitude-estimate proxy for internalised structure on a ratio scale constructed by the Chapter 53 procedure. `Chat_t` is an independently measured daily load series.

The data are **simulated**, stated plainly rather than buried: the point is to show what the suite does when the truth is known, which is impossible with real data. Generating parameters `delta = 0.15`, `eta = 0.40`, so `a_true = 0.85` and `b_true = 0.40`, process noise sd 0.06, no rating noise. **Every number below was computed, not asserted.**

The preregistration, in force before period 0

```
 period length      1 day
split           T_train = 40 (periods 0–40), T_holdout = 20 (periods 40–60)
 input variant      A
baselines       persistence; direct-input (raw u); matched AR with intercept;
                   shuffled input, three families reported
 surrogates         B = 2000 per family
 minimum u span     0.60 for the direct-input comparison to be informative
 design controls    D2 (separate channels) and D3 (prospective sampling)
 rules              Rule 1, Rule 2, Rule 3 as stated in §55.0.3
```

The training window

```
 u_t   over periods 0–39:  min 0.1248   max 0.6732   span 0.5484
 S_t   over periods 0–39:  min 0.5590   max 1.6173   mean 1.1746
                           span 1.0583 = 0.544 * ln N
 Chat_t mean over 0–39:    0.1136
 Y_t   over periods 0–40:  min 1.8722   max 2.6917   mean 2.3055   sd 0.1961
                           Y_0 = 2.0000, Y_40 = 2.3843
```

The observed span of u is 0.5484, which is below the preregistered minimum of 0.60. By S55.11 the direct-input comparison must therefore be reported as uninformative, and it will be below, even though the model happens to beat it. That is what a preregistration is for: the rule was written before the number was known, and the rule is honoured against the researcher's interest.

The fit, then the freeze

Constrained least squares on `(Y_{t+1} + Chat_t)` against `[Y_t, S_t]`, bounds a in `[0,1]`, b `>= 0`:

```
 a_hat = 0.8833      ->  delta_hat = 1 - 0.8833 = 0.1167
 b_hat = 0.3340      ->  eta_hat   = 0.3340
```

No boundary hit. Rule 3 does not fire.

Honest note on the point estimates, the most instructive line in this example. The true values are `a = 0.85` and `b = 0.40`; at `T = 40` with this noise the sampling standard deviations are 0.0341 for `b_hat` and 0.0173 for `a_hat`. The estimates sit at **-2.15 standard errors on b and +2.13 on a** — a tail draw in opposite directions, which is the correlated error the two coefficients trade between them.

A single forty-period window on a single subject recovered `eta` as 0.334 when the truth was 0.400: a 17 per cent underestimate, with everything measured perfectly and no rating noise

at all. This is why §55.11 refuses to treat one person's eta from one window as a measured constant, and it is why Chapter 56's Stage 3 asks for 120 to 240 periods rather than 40.

Frozen quantities derived from the fit, computed once and never revised:

```
  time constant      tau = 1/delta_hat = 1/0.1167 = 8.57 periods
  95% settling       ~3 tau = 25.7 periods
  implied Q_inf      (b_hat * mean S - mean Chat) / delta_hat
                   = (0.3340 * 1.1746 - 0.1136) / 0.1167 = 2.3883
                     (2.3885 at full precision)
  headroom at Y=2.39 h = eta*ln N - delta*Q - C
                   = 0.3340 * 1.945910 - 0.1167 * 2.39 - 0.115
                   = 0.6499 - 0.2789 - 0.1150 = 0.2560
```

Headroom is positive, so growth is permitted (Chapter 6). Note that `delta_hat * Q_inf = 0.2787` consumes most of the structural budget `eta_hat * ln N = 0.6499` at the fixed point — which is Theorem 3.5 doing its arithmetic, not a finding.

The held-out window: periods 40 through 60

The frozen model, scored one step ahead against observed Y_t. First eight held-out periods in full:

t	u_t	S_t	Chat_t	Y_t	predicted	actual	error
40	0.3598	1.1501	0.1079	2.3843	2.3823	2.2926	−0.0897
41	0.6597	1.6011	0.1071	2.2926	2.4526	2.3476	−0.1050
42	0.6269	1.5606	0.1280	2.3476	2.4669	2.4080	−0.0589
43	0.4145	1.2490	0.1231	2.4080	2.4210	2.5253	+0.1042
44	0.1839	0.7435	0.1003	2.5253	2.3786	2.3176	−0.0610
45	0.3606	1.1517	0.1239	2.3176	2.3078	2.2891	−0.0187
46	0.2157	0.8304	0.1210	2.2891	2.1783	2.2166	+0.0383
47	0.1662	0.6918	0.1064	2.2166	2.0826	2.0749	−0.0077

Held-out summary: mean absolute error 0.0593; held-out Y mean 2.3001, sd 0.1978; held-out u span 0.4935.

The five-row report

predictor	RMSE	skill vs model	note
frozen model	0.0777	--	a=0.8833 b=0.3340
persistence	0.1207	0.5860	
direct-input (raw u)	0.0842	0.1486	span 0.548 < 0.60: UNINFORMATIVE
matched autoregressive	0.1248	0.6123	a=0.8422, c=0.4866
shuffled input, iid	0.1267	0.6238	B=2000, p = 0.0000
shuffled input, block L=5	0.1256	0.6174	B=2000, p = 0.0000

```
shuffled input, circular     0.1300      0.6428         B=2000, p = 0.0000
```

Surrogate distributions, 5th percentile — the best that any permutation achieved:

```
iid       5th percentile 0.1162   (model 0.0777)
block     5th percentile 0.1132   (model 0.0777)
circular  5th percentile 0.1196   (model 0.0777)
```

Not one of the six thousand surrogates across the three families produced a held-out error as small as the ordered model's. The best single permutation, out of 2000 block permutations, still predicted 46 per cent worse than the true ordering.

The three rules, applied

```
RULE 1  SHUFFLE      Calibrated critical value from the preregistered design
                  simulation: reject at p <= 0.05 for the block family, whose
                measured false-positive rate on this design was 0.060 against
                       a nominal 0.050. Observed p = 0.0000 with B = 2000.
                       DOES NOT FIRE.

RULE 2  AUTOREGRESSION  Skill against matched AR = 0.6123, positive and large.
                        DOES NOT FIRE.

RULE 3  CONSTRAINTS  a_hat = 0.8833 strictly inside [0,1]; b_hat = 0.3340 > 0;
                     no boundary hit; unconstrained fit also inside the box.
                        DOES NOT FIRE.
```

Verdict: the model survives its own falsification suite on this subject and this window, with `delta_hat = 0.117 +- 0.017` and `eta_hat = 0.334 +- 0.034` at one standard error, both point estimates biased by roughly two standard errors in this particular draw. That last clause belongs in the abstract, not the appendix.

The same suite, run on a null

Now the honest negative, the half of this chapter that matters most. Same procedure, same period length, same instrument, same forty training periods — but Y is generated by a smooth process that has nothing to do with S. There is no relationship to find. The suite should say so.

```
fitted a_hat = 1.0000  ->  delta_hat = 0.0000    BOUNDARY HIT: Rule 3 event
fitted b_hat = 0.0931  ->  eta_hat   = 0.0931

in-sample R^2 on the training window: 0.6477

held-out RMSE   frozen model            0.0978
                persistence             0.0730
                direct-input            0.1116
                matched autoregressive  0.0681
              shuffled iid, mean     0.0793    p(shuffle <= model) = 0.999
              shuffled circular, mean 0.0763   p(shuffle <= model) = 1.000
```

```
skill vs matched autoregressive: −1.0669
```

Read that block slowly; it is the shape of a false result and every line is diagnostic.

The in-sample R^2 is 0.6477, which presented alone reads as a solid finding. The frozen model is then beaten on held-out data by persistence (0.0730 against 0.0978), beaten decisively by matched autoregression (0.0681, skill -1.0669), and beaten by essentially every one of 2000 circular-shift surrogates — p = 1.000, meaning the permuted inputs predicted *better* than the true ordering, which is what happens when the true ordering carries no information. And the constrained fit hit the wall: `a_hat = 1.0000` exactly, the optimiser pressing against a boundary it was not allowed to cross. By S55.6 that goes in the abstract.

```
RULE 1  SHUFFLE        p = 1.000.  FIRES.
RULE 2  AUTOREGRESSION skill = −1.0669.  FIRES.
RULE 3  CONSTRAINTS    a_hat at the boundary.  FIRES.

VERDICT: FALSIFIED on this subject and this window. All three rules.
```

The correct report, in full:

> *We preregistered the minimal estimation model of Paper E §5.2 for a single subject over 61 daily periods, fitted on 40 and froze. The frozen model failed on held-out data against all four baselines. The shuffled-input control, run with B = 2000 circular-shift surrogates and a critical value calibrated by simulation on our own design, predicted held-out order as well as or better than the true input ordering (p = 1.000). The matched autoregressive baseline beat the model with a skill score of -1.0669. The constrained fit hit the retention boundary at a = 1.000. All three preregistered falsification rules fired. In this subject, over this window, with this operationalisation, internal surplus does not predict stored order.*

That paragraph is a good outcome. It cost the same as the other one, it is fully informative, and it is publishable under P8. **The suite that produced the confirmation in the first half of this example is the same suite, unchanged, that produced the refutation in the second half, and neither result required the analyst to exercise a single judgement after the freeze.**

What this example does not establish

The limits, because the example is seductive:

```
It does not establish that the model holds for this subject in general.
    One window. Chapter 56 Stage 3 asks for two, then more.

It does not establish that eta = 0.334 is a property of this person.
    The measured noise floor at T=40 is far above the difference between
    0.334 and 0.400. FVP-7 has not been attempted here.

It does not establish that the logarithmic input form is correct.
    The u span was 0.548, below the preregistered minimum. Reported.

It does not rule out fabricated timing.
    D2 and D3 were in force; D1 (blinding) was not. Under uncontrolled
```

```
        conditions the measured pass rate for a fabricated relationship is
        52.7 per cent, so this omission is material and is stated as such.

    It does not establish anything at group scope.
        Chapter 56 Stage 5.
```

Five limitations, each named, each with the study that would remove it. That list is not an apology appended to a result. **It is the result: a claim, its exact boundary, and the next experiment.**

A framework that cannot fail cannot be trusted. This one fails at three named rules, on data whose truth we control, in the second half of its own worked example — and the next chapter lays out, in order and with its kill conditions, the programme that would put those rules in front of real people.

Chapter 56 . The Research Program

The Law

This chapter proves nothing and measures nothing. It states a plan.

Badged as the author's programme, not as a result: the five stages below, run in order with the kill conditions stated at each, would settle whether the formulas of Volume Two are science. Every number in it is either arithmetic from a declared assumption or a simulation of a declared design. No stage has been run. Saying so in the first paragraph is the only honest way to open a chapter of this kind.

```
THE PROGRAMME

  STAGE 1   PILOT OPERATIONALISATION
            Fix what D, P, and E mean for one subject, in units.
            Gate: FVP-1 (dimensions), FVP-3 (limiting cases)

  STAGE 2   RATIO-SCALE CONSTRUCTION
            Build ratio scales for D, P, E by magnitude estimation and
            cross-modality matching. Verify transitivity.
            Gate: FVP-2 (scale type) -- the gate that licenses division

  STAGE 3   SINGLE-SUBJECT REPEATED MEASURES
            Estimate delta and eta within one person, repeatedly, and ask
            whether they are stable properties of that person.
            Gate: FVP-5 (identifiability), FVP-6 (falsification)

  STAGE 4   CROSS-SUBJECT PARAMETER RECOVERY
            Same operationalisation, many people, two labs. Do the
            parameters recover?
            Gate: FVP-7 (reproducibility) -- MAKE OR BREAK

  STAGE 5   PREREGISTERED FIELD TRIALS AT GROUP SCOPE
            Cluster-randomised, time-series, held-out prediction.
            Gate: FVP-6 at the external scope; the correspondence claim of
            Chapter 50 becomes testable here and only here
```

The ordering is not a preference. **Each stage consumes the output of the one before it, and running any stage before its predecessor has cleared produces a number that cannot be interpreted.** Estimating `eta` (Stage 3) on a scale not shown to be ratio-scaled (Stage 2) yields a coefficient in unknown units; comparing it across people (Stage 4) compares unknown units to unknown units; deploying it at group scope (Stage 5) does so at institutional expense.

56.0.1 The governing principle

A program that cannot be killed at any stage is not a research program. It is a research aspiration with a schedule attached.

Every stage below carries four blocks, and the third is what makes it a program:

```
WHAT IS MEASURED
WHAT WOULD ADVANCE IT
WHAT WOULD KILL IT AT THIS STAGE
SCALE AND TIMESCALE
```

The kill conditions are written to be recognisable in advance and unpleasant to reach. A kill condition a determined investigator can argue past is a decoration; §56.10 gives the tripwires for the specific ways such arguments get made.

56.0.2 What is proven and what is postulated

Proven. Nothing empirical. The mathematics used throughout — Paper A Theorem 2.1, Paper D Corollary 2.2, Paper E §3 and §5 — is proven and is being *applied*, not tested. **The theorems cannot fail; what can fail is the claim that the world's psychological and social quantities are the objects those theorems are about.** That is why a correct mathematical result is no evidence at all that the program will succeed.

Also proven, used here as design arithmetic and not as findings: Spearman-Brown, the Fisher-z interval for a correlation, the standard error of a log-log slope, the cluster design-effect formula, and the exact scaling `b_hat -> b_hat/g` under an input inflated by `g` (Chapter 55 §55.11). All algebra, computed below, not cited.

Postulated, badged. That `D`, `P`, and `E` can be operationalised at all; that the storage law's form applies internally (CANON_INTERNAL §8); that `delta` and `eta` are person-level constants over months rather than moment-level variables; that Chapter 50's correspondence is more than an analogy. **Every one is exactly what a stage below is designed to break.**

Honesty clauses, repeated. The triangle inequality fails for geometric triples (Paper A §5); sub-additivity holds on u only. As deployed, the surplus functional is a lower bound on the ensemble's information — we log the floor, not the ceiling. Paper A is pure mathematics and makes no empirical or metaphysical claim. Paper D does not identify surplus with Shannon entropy or mutual information. Paper E proves no thermodynamic negentropy law: `F(u) + N_can(u) <= ln N` (Corollary 2.2), and higher static surplus does not imply higher canonical concentration.

56.0.3 The field this program is entering

Psychology and sociology are not pseudoscience. They are sciences with a measurement problem and a replication problem. The Open Science Collaboration (*Science*, 2015) attempted 100 replications: 97 originals had reported significant results, about 36 per cent of replications did, effect sizes roughly halved.

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does. Psychophysics already achieves it — Weber (1834), Fechner (1860), Stevens (1957) — and psychometrics comes close through the Rasch model's separation of person and item parameters (Chapter 54). This program is an attempt to bring Volume Two's formulas onto that footing, not an attempt to prove them.

The Engineering

56.1 Stage 0 — preconditions

Four things must exist before Stage 1. None is research; starting without them guarantees a Stage 1 that cannot be interpreted.

```
P1  A WRITTEN THEORY OF THE SUBJECT.
    Which construct is Q for this subject? Not "wellbeing". Something
    that could in principle be destroyed by a year of neglect and rebuilt
    by a year of practice: a skill, a body of knowledge, a settled habit
    of attention, a repertoire.

P2  A CANDIDATE ORDER PROXY THAT IS NOT A REPACKAGING OF U.
    Paper E §6.3. Write the argument down and test it: could Y rise in a
    period of intense interaction that produced nothing durable? If no,
    Y is measuring U.

P3  AN INDEPENDENT COST INSTRUMENT.
    C in the same units as eta*S. If cost cannot be measured, the model
    is fitted without it and is a weaker model (Chapter 55, F55.1).

P4  ETHICS AND CONSENT INFRASTRUCTURE, INCLUDING COSTLESS REFUSAL.
    CANON §5, refusal 4: no intervention removes a person's ability to
    decline. In a measurement programme that means refusal to be measured
    must be available, costless, and recorded as ordinary missing data.
```

Stage 0 has no kill condition because it has no result. It has an entry condition: all four, in writing, before recruitment.

56.2 Stage 1 — Pilot operationalisation of D, P, and E for a single subject

What is measured One subject, 45 sessions, three candidate instruments per term:

```
D  PERSPECTIVE, the frame brought to the moment.
     D-a  prior-expectation elicitation, with a confidence, before the
          episode
     D-b  frame inventory: which of the subject's N audited framings
          (Chapter 39) were active, weighted
     D-c  magnitude estimate of "how strongly I already knew what this
          would be like", against a fixed modulus

P  PERCEPTION, the intake from the moment.
     P-a  detail-recall count, scored against a record
     P-b  discrepancy report: what departed from D-a, with magnitude
     P-c  magnitude estimate of "how much came in", same modulus

E  ENVIRONMENT, the situation itself.
     E-a  independently coded situational richness from a record
     E-b  physical-parameter panel where the room is the environment
```

```
          (Chapter 12: acoustics, light, sittable space, sightlines)
    E-c   magnitude estimate of "how much was there to be had"
```

The purpose is not to find the right instrument. **It is to find out whether `D` and `P` can be put in the same units at all**, because FVP-1 requires `r = D/P` to be dimensionless — a far stronger demand than that both be numbers. This is the least glamorous stage and the one most often skipped by starting straight from a questionnaire.

Per candidate: session-level reliability (internal consistency across items), test-retest reliability (across sessions, short interval), and convergent correlation (-a against -b against -c).

What would advance it

```
A1  One operationalisation of D and one of P share units, with a written
    argument for the shared unit and a demonstrated common zero.

A2  Session-level reliability >= 0.80 for at least one instrument per term.

A3  Test-retest r >= 0.70 across a one-week interval, 95 per cent lower
    bound above 0.50.

A4  FVP-3's four limiting cases are observable and recognisable: P -> 0
    (frame with no intake), D -> 0 (intake with no frame), E -> 0, and
    D = P (calibrated contact). If none can be produced or recognised in
    a real session, the formula's structure has no referent.
```

A3 sets the sample size. The 95 per cent interval on an observed correlation, by Fisher-z:

n	95% CI on an observed r = 0.70	width
10	[0.126, 0.923]	0.797
20	[0.373, 0.872]	0.499
30	[0.454, 0.847]	0.392
50	[0.524, 0.819]	0.295
80	[0.568, 0.797]	0.230
120	[0.595, 0.781]	0.186
200	[0.622, 0.764]	0.143

The smallest `n` at which an observed `r = 0.70` has a 95 per cent lower bound above 0.50 is 41. Thirty sessions is therefore under-powered for A3, and the honest choices are 41 or more sessions, or relaxing A3 to a lower bound above 0.40 and saying so in advance. This manual takes the first and sets Stage 1 at 45 sessions.

A2 sets the item count. Spearman-Brown, `r_k = k*r1/(1 + (k-1)*r1)`:

single-item r1	items for reliability 0.80	for 0.90
0.10	37	82
0.15	23	52
0.20	17	37
0.30	10	22
0.40	7	14
0.50	5	10

If the single-item reliability of a P instrument is 0.15, a session score needs 23 items to reach 0.80 and 52 to reach 0.90. That is what decides session length, and why Stage 1's session is 25 minutes rather than five.

What would kill the program at this stage

```
K1.1  NO SHARED UNIT. After honest effort, no operationalisation of D and
      P can be expressed in the same units with a common zero.
      CONSEQUENCE: D/P is not a ratio. M = (D/P)E is not an equation.
      By FVP-2 it must be demoted in print to a heuristic ordering, and
      the Moment equation ceases to be a formula in this book's sense.
      The storage law survives; the Moment equation does not.

K1.2  NO RELIABLE INSTRUMENT. No candidate reaches session reliability
      0.80 with a session under 45 minutes.
      CONSEQUENCE: Chapter 55 §55.9's attenuation applies with force -- at
      rating sd 0.40 against signal sd near 0.20, delta was estimated at
      0.432 against a true 0.150. Estimates on such an instrument are not
      noisy but systematically wrong. The programme cannot proceed.

K1.3  THE LIMITING CASES HAVE NO REFERENT. Neither the subject nor a
      trained observer can identify instances of P -> 0, D -> 0, or
      D = P in the subject's own record.
      CONSEQUENCE: FVP-3 fails. A formula whose divergences correspond
      to nothing nameable is not describing the phenomenon.

K1.4  Y IS A REPACKAGING OF U. The candidate order proxy rises with
      interaction volume and with nothing else.
      CONSEQUENCE: Paper E §6.3's named error. Return to Stage 0 P2 or
      stop.
```

K1.1 is the most likely kill in the whole program and the one people will most want to argue past. **That "perspective" and "perception" are expressible in a common unit is not obvious, is not established, and is the load-bearing assumption of the entire Moment equation.** The honest outcome of a failed Stage 1 is a chapter saying the Moment equation is a heuristic ordering — useful for diagnosis, illegitimate for arithmetic — and Chapter 53 already specifies what that demotion looks like in print.

Scale and timescale

```
subjects            1  (plus 1 replication subject, run in parallel)
sessions           45  at 3 per week
session length     25 minutes x 3 instruments per term
subject burden     45 x 25 x 3 / 60 = 56.3 subject-hours
fieldwork          45 / 3 = 15 weeks = 3.5 months
total stage        7 months including instrument development and analysis

researcher hours (itemised, assumptions declared not cited):
```

```
instrument development and piloting          160.0
session supervision                           37.5
coding at 2x subject time                     75.0
analysis and FVP-1/2/3 write-up               80.0
preregistration drafting                      24.0
------------------------------------------------
TOTAL                                        376.5
```

Seven months and about 377 researcher-hours to find out whether the central formula of Part VII has units. A cheap way to learn something decisive, and the correct first expenditure because a failure here saves the other five years.

56.3 Stage 2 — Ratio-scale construction

What is measured Magnitude estimation and cross-modality matching on the instruments that survived Stage 1.

```
MAGNITUDE ESTIMATION
  A modulus is presented and assigned a number; each subsequent stimulus
  is assigned a number in proportion. Fit log psi = log k + a log I and
  read the exponent a.

CROSS-MODALITY MATCHING
  The subject adjusts a stimulus in modality B until it matches one in
  modality A. The matching exponent is predicted by the ratio of the two
  modalities' own exponents. This is the transitivity test, and it is the
  actual evidence that the scale is a ratio scale and not a numbering
  habit.
```

Published, replicated exponents, nominal, and constants of a *procedure* rather than of nature — they vary with method, range and instruction:

```
Loudness (1 kHz)     a ~ 0.67     compressive
Brightness           a ~ 0.33     compressive
Visual length        a ~ 1.0      linear
Taste (salt)         a ~ 1.3      expansive
Heaviness (lifted)   a ~ 1.45     expansive
Warmth on arm        a ~ 1.6      expansive
Electric shock       a ~ 3.5      strongly expansive
```

These are the calibration ladder. **D, P, and E are anchored by matching against modalities whose exponents are already known, which is what converts a number the subject produced into a measurement.**

What would advance it

```
B1  Each of D, P, E yields a stable within-subject exponent across
    sessions, with a usefully small standard error.

B2  Cross-modality matching is TRANSITIVE within measurement error:
```

```
    a_AC = a_AB * a_BC.

B3  D and P yield exponents consistent with a common unit, so r = D/P is
    dimensionless (Stage 1's A1, now quantitative).

B4  A true zero is demonstrable for both D and P: a condition reliably
    assigned magnitude zero, not merely the smallest number used.
```

The precision arithmetic sets the number of judgments. For a log-log regression, `SE(a) = sigma_resid / (sd(log I) * sqrt(n))`, and with stimulus levels spread evenly over a span of L decades, `sd(log I) = L/sqrt(12)`:

```
residual sd 0.25 log10 units; stimulus span in decades:

  n \ span     1.0      1.5      2.0      3.0
    20        0.1936   0.1291   0.0968   0.0645
    40        0.1369   0.0913   0.0685   0.0456
    60        0.1118   0.0745   0.0559   0.0373
   100        0.0866   0.0577   0.0433   0.0289
   200        0.0612   0.0408   0.0306   0.0204
   400        0.0433   0.0289   0.0217   0.0144

judgments needed at a 2.0-decade span:
   SE(a) <= 0.05  ->  n =  75
   SE(a) <= 0.03  ->  n = 209
   SE(a) <= 0.02  ->  n = 469
```

Two design decisions fall out of that table. First, span matters as much as count: going from 1.0 to 2.0 decades halves the standard error at every n, the same gain as quadrupling the judgments. Widen the range before lengthening the session. Second, `SE(a) <= 0.02` costs 469 judgments per exponent — a real afternoon per term per subject.

The transitivity test's power depends on that precision. Take a chain with `a_AB = 0.67` and `a_BC = 1.45`, so predicted `a_AC = 0.9715`:

```
 SE per exponent   predicted a_AC     SE of (a_AC - a_AB*a_BC)   detectable at 95%
      0.02          0.9715 +- 0.0319          0.0377                  0.0739
      0.05          0.9715 +- 0.0799          0.0942                  0.1847
      0.10          0.9715 +- 0.1597          0.1885                  0.3694
```

At `SE = 0.10` per exponent, a transitivity violation smaller than 0.37 cannot be detected, and a violation of 0.37 is enormous. A transitivity test on imprecise exponents is not a test but a guarantee of agreement. This is how the ratio-scale claim gets made without being earned, and why B1's precision requirement precedes B2's transitivity requirement.

Session arithmetic for the stage:

```
9 stimulus levels x 6 replications x 6 sessions x 3 modalities
    = 972 judgments per subject per term
x 3 terms (D, P, E)
    = 2,916 judgments per subject
at 8 seconds per judgment
```

```
    = 6.48 subject-hours per subject
6 subjects
    = 38.9 subject-hours total
```

972 judgments per term per subject clears the `n = 469` requirement for `SE(a) <= 0.02` on each modality separately, which is the precision the transitivity test needs.

What would kill the program at this stage

```
K2.1  NO STABLE EXPONENT. The fitted exponent for D or P varies across
      sessions by more than its own standard error by a wide margin.
      CONSEQUENCE: the subject is not using the number line consistently.
      There is no scale to build on. FVP-2 fails.

K2.2  TRANSITIVITY VIOLATED. a_AC differs from a_AB * a_BC by more than
      the detectable threshold above.
      CONSEQUENCE: the judgments are not measurements of a single
      underlying magnitude. Classic and decisive evidence that a putative
      ratio scale is not one.

K2.3  NO TRUE ZERO. No condition is reliably assigned magnitude zero for
      D or for P.
      CONSEQUENCE: the scale is at best interval, and by Stevens'
      typology division is impermissible on an interval scale. D/P is
      meaningless as written. Demote the Moment equation in print.

K2.4  D AND P LAND IN INCOMPATIBLE UNITS. Their exponents cannot be
      reconciled to a common unit.
      CONSEQUENCE: as K1.1, now with numbers attached.
```

K2.3 is worth stating twice. **A subjective quantity with no true zero cannot be divided by another, and the entire content of `M = (D/P)E` is a division.** The program's most probable non-trivial outcome is not that the Moment equation is false but that it is not an equation — that `r = D/P` survives as an ordering (Chapter 35) while ceasing to be a number. That is a genuine finding, and Chapter 53 specifies how to print it.

Scale and timescale

```
subjects            6
judgments       2,916 per subject; 17,496 total
subject burden   6.48 hours per subject; 38.9 hours total
fieldwork        6 sessions per term at 1 per week; terms run in parallel
                 = 6 weeks
total stage      4 months including apparatus construction and analysis

researcher hours:
   session supervision at 1.0x subject time         38.9
   stimulus and apparatus construction, 3 modalities  240.0
   scale fitting and transitivity analysis         180.0
```

```
independent replication wave                    200.0
writing                                         120.0
-------------------------------------------------------
TOTAL                                           778.9
```

56.4 Stage 3 — Single-subject repeated measures

What is measured Three subjects who cleared Stages 1 and 2, sampled four times daily for 240 days. Per period: `Y_t` on the constructed scale, `U_t` over the subject's own audited framings, `Chat_t` from the independent cost instrument. The Chapter 55 suite is applied per subject: fit on the first 120 periods, freeze, predict the second 120, four baselines, three rules.

The question is not "does the model fit" but **"are `delta` and `eta` properties of this person, or of the window in which they were estimated?"**

What would advance it

```
C1  The frozen model clears all three falsification rules on held-out
    data, per subject, with a calibrated surrogate test.

C2  Within-person split-half reliability of b_hat is at least 0.80.

C3  a_hat lies strictly inside [0,1] in every subject, no boundary hit.

C4  No systematic drift: first-half parameters predict the second half,
    and second-half parameters predict the first, to within sampling
    error.
```

The precision arithmetic, simulated at true `a = 0.85`, true `b = 0.40`, process noise 0.06, rating noise 0.10, 400 replications per row:

T	mean a_hat	sd a_hat	mean b_hat	sd b_hat	95% CI half-width on b
30	0.8174	0.0409	0.4604	0.0781	0.1532
60	0.8224	0.0248	0.4521	0.0494	0.0968
90	0.8242	0.0207	0.4487	0.0411	0.0805
120	0.8236	0.0166	0.4513	0.0331	0.0649
180	0.8252	0.0137	0.4485	0.0278	0.0545
240	0.8248	0.0119	0.4488	0.0243	0.0476
360	0.8257	0.0097	0.4476	0.0198	0.0388

Two features of that table are the reason Stage 3 exists.

First, the bias. At every T, `b_hat` sits near 0.449 against a true 0.400 — a persistent 12 per cent upward bias that does not shrink with more data, because it is not sampling error but the attenuation of §55.9 at rating noise 0.10. More periods buy precision, not accuracy. **Only a better instrument buys accuracy, and that is Stage 2's job.**

Second, the cost of precision. Halving `sd(b_hat)` from 0.0494 to 0.0243 means going from 60 periods to 240 — four times the fieldwork for twice the precision, the ordinary square-root law, worth stating in weeks:

```
T = 120 at daily sampling      = 17.1 weeks =  3.9 months
T = 120 at twice-weekly        = 60.0 weeks = 13.8 months
T = 240 at daily sampling      = 34.3 weeks =  7.9 months
```

The middle row is why daily sampling is not optional. **Twice-weekly sampling turns a four-month study into a fourteen-month study for the same precision, and a fourteen-month study has a different `delta` at the end than at the beginning.**

Power to detect a genuine within-person *change* in `eta` between two blocks, with `SE_diff = sqrt(2) * sd_b`:

T	SE_diff	power at d=0.05	d=0.10	d=0.20
60	0.0698	0.110	0.299	0.817
90	0.0581	0.138	0.406	0.931
120	0.0469	0.187	0.569	0.990
180	0.0393	0.246	0.720	0.999
240	0.0344	0.307	0.829	1.000
360	0.0280	0.432	0.947	1.000

Two blocks of 120 periods give 57 per cent power to detect a change in `eta` of 0.10 and 19 per cent for a change of 0.05. Any claim that a practice, therapy, habit or season changed a person's conversion efficiency by a small amount is, at these sizes, undetectable. Stating that in advance prevents the commoner outcome, which is detecting it anyway.

Burden: 960 prompts per subject at about 90 seconds each — 24 subject-hours over eight months. Adherence is itself a threat; it must be measured and reported, and missing periods handled by the preregistered rule, never by interpolation chosen after the fact.

What would kill the program at this stage

K3.1 ANY OF THE THREE RULES FIRES IN A MAJORITY OF SUBJECTS.
Shuffle not beaten, matched AR not beaten, or a boundary hit.
CONSEQUENCE: the storage law does not describe this subject's internal dynamics. Report it (Chapter 55, P8) and stop.

K3.2 SPLIT-HALF RELIABILITY OF b_hat BELOW 0.50.
The simulated benchmark at 60 periods per half is 0.841 and at 240 per half is 0.962. A real figure below 0.50 means the parameter is a property of the window.
CONSEQUENCE: eta is not a person-level constant on the months timescale. Stage 4 is meaningless and must not be attempted.

K3.3 SYSTEMATIC SAME-DIRECTION DRIFT in a_hat or b_hat across blocks, in every subject.
CONSEQUENCE: the constant-parameter postulate fails. A time-varying rescue must be preregistered as a NEW prediction (§56.10), never adopted as an explanation of this failure.

K3.4 THE DESIGN CONTROLS FAIL. Debriefing reveals subjects inferred the hypothesis, and no blinding or channel separation was in force.

```
CONSEQUENCE: Chapter 55 §55.8 measured a 52.7 per cent pass rate
for fabricated relationships under those conditions. The data
cannot be interpreted however good they look.
```

K3.4 will be tempting to waive because the data will already exist. **Waiving it converts eight months of fieldwork into eight months of fieldwork that means nothing — exactly the outcome the waiver is trying to avoid and exactly the outcome it produces.**

Scale and timescale

```
subjects           3  (each analysed separately; no pooling at this stage)
periods          240  per subject, daily, 4 prompts per day
prompts          960 per subject; 2,880 total
subject burden    24 hours per subject; 72 hours total
fieldwork       34.3 weeks = 7.9 months
total stage       12 months including pipeline build and analysis

researcher hours:
   adherence monitoring (3 x 34.3 weeks x 0.5 h)     51.4
   data pipeline build                              200.0
   surrogate calibration simulations                120.0
   per-subject fits, split-half, analysis           300.0
   writing                                          160.0
   ------------------------------------------------------
   TOTAL                                            831.5
```

56.5 Stage 4 — Cross-subject parameter recovery: the make-or-break stage

What is measured Forty subjects, 120 daily periods each. **The identical protocol run independently in two laboratories that have never met, working from the published Stage 1-3 protocol alone.**

That last clause is the stage. Everything before it can be done by one group with one set of habits. FVP-7 asks whether the parameters survive a different group's reading of the same written procedure, and CANON_INTERNAL §6 names it the gate that decides the hard-science question and the one most often skipped. Three questions:

```
Q1  Do individual differences in delta and eta EXIST -- is the spread
    across people larger than the estimator's noise?

Q2  Do they RECOVER -- does a person's estimated parameter correlate with
    that person's true parameter, insofar as a second estimate can serve
    as a proxy for truth?

Q3  Do they TRANSPORT -- does lab 2 get the same numbers as lab 1, and if
    not, is the disagreement a single scaling constant or a pattern?
```

The noise floor, and why Q1 comes first The arithmetic most individual-differences research omits. Simulate sixty subjects who share one true `eta` and see how much variation the estimator

invents:

```
THE NOISE FLOOR -- 60 subjects who truly share eta = 0.40

  T per subject      sd of b_hat across subjects
      60                    0.0478
     120                    0.0355
     240                    0.0255
     360                    0.0166
```

At 120 periods per subject, an estimator applied to sixty identical people reports a between-person standard deviation of 0.036 in `eta`. Any individual-differences claim resting on a smaller spread is a claim about the estimator, not about people. An independent 30-subject run with the true spread set to zero produced 0.042: same conclusion, different seed.

The consequence for Q2, with 30 subjects at 120 periods:

```
true between-person sd of eta       corr(true, estimated)
           0.12                            0.956
           0.04                            0.503
```

The second row is not an estimator failure. **It is what happens when the quantity varies less across people than the instrument's error, and no number of subjects repairs it — only longer series per subject or a quieter instrument.**

Power to reject "everyone shares one `eta`", testing the observed spread against a simulated noise floor at one-sided 5 per cent:

```
  n     T      true between-sd     critical sd     power
 20   120          0.05              0.0415         0.97
 20   120          0.10              0.0428         1.00
 40   120          0.05              0.0391         0.98
 40   120          0.10              0.0407         1.00
 20   240          0.05              0.0298         1.00
 40   240          0.05              0.0282         1.00
```

Twenty subjects at 120 periods already gives 97 per cent power to detect a true between-person standard deviation of 0.05 in `eta`. Forty is comfortable, and the binding constraint is periods per subject, not subjects — the opposite of the usual instinct. **One more month of daily sampling per subject buys more than ten more subjects.**

Reliability of the person-level estimate, and the ceiling it puts on everything downstream:

```
  T     between-sd    var_err     reliability    max corr with any
                                                 external variable
120       0.05       0.00126        0.665            0.815
120       0.10       0.00126        0.888            0.942
120       0.15       0.00126        0.947            0.973
240       0.05       0.00065        0.793            0.891
240       0.10       0.00065        0.939            0.969
240       0.15       0.00065        0.972            0.986
```

Before correlating `eta` with anything, compute this ceiling and publish it: an observed correlation cannot exceed it, and one reported near it is evidence of a problem, not a discovery.

Q3, and the sharpest criterion in the program Two labs, same protocol, different hands. Lab 2's operationalisation of u runs 15 per cent hot: a wider framing set, a different inner product, an entirely defensible reading of the procedure. Simulated:

```
lab 1 mean b_hat        0.4679
lab 2 mean b_hat        0.4068
observed ratio          0.8696
predicted ratio 1/1.15  0.8696
```

Exact, because the model is linear in its input: scaling S by g scales `b_hat` by `1/g`. Other calibration errors, equally exact:

```
inflation g      predicted b_hat ratio 1/g
   1.05                 0.9524
   1.10                 0.9091
   1.15                 0.8696
   1.25                 0.8000
   1.50                 0.6667
```

The criterion:

```
BETWEEN-LAB CRITERION

A CONSTANT FACTOR across all subjects, matching 1/g for an independently
measurable operationalisation difference g
    -> CALIBRATION difference. Repairable. Publish the constant.

A FACTOR THAT VARIES ACROSS SUBJECTS, or any change in the RANK ORDER of
subjects between labs
    -> THEORY FAILURE. The parameter is not a property of the person.
```

A framework passes FVP-7 when between-lab disagreement is a single number, and fails when it is a pattern. That is why Stage 4 is make-or-break: it is the only stage whose failure cannot be attributed to an instrument, a subject, or a window.

What would advance it

```
D1  The between-person spread in eta exceeds the simulated noise floor at
    the study's own n and T, by a margin, in both labs.

D2  Rank order of subjects on eta agrees between labs, to a correlation
    significantly above zero and at or below the published reliability
    ceiling.

D3  Any mean-level disagreement is a single scaling constant matching an
    independently measured operationalisation difference.
```

```
D4  Both labs clear the Chapter 55 three-rule suite in most subjects.
```

What would kill the program at this stage

```
K4.1  NO INDIVIDUAL DIFFERENCES. The between-person spread does not clear
      the noise floor at n = 40, T = 120, in either lab.
      CONSEQUENCE: eta and delta are not person parameters. Chapter 45's
      Personal Formulary describes a species-typical process, not
      individual constants, and every "your delta" in Volume Two must be
      rewritten. Survivable, but a large retraction, printed as one.

K4.2  NO RANK-ORDER AGREEMENT BETWEEN LABS. Lab 1 and lab 2 order the same
      forty subjects differently on eta, beyond sampling error.
      CONSEQUENCE: FVP-7 fails outright. The parameters are properties of
      a procedure and not of a person. THE PROGRAMME STOPS. There is no
      rescue that preserves the claim, because the claim was that these
      are measurements.

K4.3  MEAN DISAGREEMENT THAT IS NOT A CONSTANT FACTOR. Labs differ in a
      way that varies across subjects and does not track any measurable
      difference in operationalisation.
      CONSEQUENCE: as K4.2. An unexplained, subject-varying between-lab
      discrepancy is the signature of a parameter that does not exist
      independently of who measured it.

K4.4  THE SUITE FAILS IN ONE LAB AND PASSES IN THE OTHER, with no
      identifiable procedural difference.
      CONSEQUENCE: the strongest possible warning. Do not proceed to
      Stage 5. Return to Stage 1 and find the procedural difference, or
      stop.
```

K4.2 is the program's true kill switch. **If the same forty people get different `eta` orderings from two competent labs following the same written protocol, then `eta` is not a quantity the people have; it is a quantity the procedure makes, and no further work changes that.** Chapter 54's specific objectivity is the standard being failed: a person's measure depending on which instrument was used is precisely what a measurement is not allowed to do.

Scale and timescale

```
subjects           40 per lab; 80 total across two labs
periods           120 per subject, daily, 4 prompts per day
prompts           480 per subject; 38,400 total
subject burden     12 hours per subject; 960 hours total
fieldwork        17.1 weeks = 3.9 months per lab cohort
total stage        18 months including recruitment, two cohorts,
                   analysis and between-lab reconciliation

researcher hours:
```

```
fieldwork, two labs (recruitment, monitoring, coding)   822.0
data pipeline adaptation                                 150.0
analysis, noise floor, between-lab reconciliation        500.0
writing                                                  240.0
------------------------------------------------------------
TOTAL                                                   1712.0
```

56.6 Stage 5 — Preregistered field trials at group scope

What is measured Only now does the program touch institutions, and only now is it entitled to. Two tests, which must not be confused with each other:

```
T1  Does the storage law predict held-out group order, at group scope,
    against the same four baselines? (Chapter 55, external scope)

T2  Does an intervention chosen by the framework -- raise eta, cut C, or
    raise u, whichever the headroom gauge says is cheapest (Chapter 26)
    -- outperform an active control chosen by ordinary practice?

T1 is a prediction test. T2 is a causal test. A design that runs only T2
can succeed while the framework is wrong (the intervention worked for
other reasons) and can fail while the framework is right (the
intervention was implemented badly). BOTH ARE REQUIRED.
```

And a third, the one that justifies Volume Two's existence:

```
T3  Does the internal/external correspondence of Chapter 50 hold
    QUANTITATIVELY? Are members' internal eta values, measured by Stage
    3-4 methods, related to their group's eta? Chapter 50 asks for the
    cases where the correspondence BREAKS DOWN to be stated honestly, and
    Stage 5 is where they would be found.
```

Design: cluster-randomised, 24 clusters, 20 members each, 52 weekly periods, framework intervention in half and an active control in the other half. Time series throughout, not endpoints.

The sample-size arithmetic `z(0.975) = 1.959964`, `z(0.80) = 0.841621`, so `(z_a + z_b)^2 = 7.848880`. Individual-level n per arm for 80 per cent power at two-sided alpha = 0.05, from `n = 2*(z_a+z_b)^2 / d^2`:

```
d = 0.20      n per arm = 393
d = 0.30      n per arm = 175
d = 0.50      n per arm =  63
d = 0.80      n per arm =  25
```

Clustering inflates this by `DE = 1 + (m-1)*ICC`:

m \ ICC	0.01	0.02	0.05	0.10	0.20
5	1.04	1.08	1.20	1.40	1.80
10	1.09	1.18	1.45	1.90	2.80
20	1.19	1.38	1.95	2.90	4.80
30	1.29	1.58	2.45	3.90	6.80

50	1.49	1.98	3.45	5.90	10.80
100	1.99	2.98	5.95	10.90	20.80

Clusters per arm at d = 0.30, 80 per cent power (individual n per arm = 175):

m	ICC=0.02	ICC=0.05	ICC=0.10
8	25	30	38
12	18	23	31
20	13	18	26
30	10	15	23
50	7	13	21

At 20 members per cluster and ICC 0.05, eighteen clusters per arm — 36 in total — are needed for 80 per cent power on a standardised effect of 0.30. The design above specifies 24, which is under-powered for d = 0.30 at ICC = 0.05 and adequate for d = 0.50. Stated here rather than discovered later: **24 clusters is a pilot-scale trial for a moderate effect and a full-scale trial only for a large one, and the preregistration must say which claim it is entitled to make.**

Time-series length, since T1 requires held-out prediction and not a pre-post comparison:

periods	cadence	weeks	years	train / hold out
26	fortnightly	52	1.00	17 / 9
52	fortnightly	104	2.00	34 / 18
52	weekly	52	1.00	34 / 18
104	weekly	104	2.00	69 / 35

Fifty-two weekly periods with a 34/18 split is the minimum giving an 18-period held-out window, which the Chapter 55 surrogate test needs for resolution at B = 999.

What would advance it

E1 T1: the frozen model beats all four baselines on held-out group order, with a calibrated surrogate test, in most clusters.

E2 T2: the framework-selected intervention outperforms the active control on the preregistered primary outcome.

E3 The headroom gauge PREDICTED which clusters would fail to grow. Corollary 3.4 says that at h_t < 0 no interaction geometry prevents decline, so clusters with negative measured headroom should not have grown whatever was done to them. The sharpest external prediction the framework makes.

E4 T3: a measurable relation between members' internal parameters and the group's, with the breakdown cases identified and printed.

E3 is the strongest test available at group scope. **Corollary 3.4 is a prohibition, and prohibitions are the most falsifiable things a theory can offer: a single cluster with clearly negative headroom that nonetheless grew its order proxy substantially, with eta, C, and N unchanged, is a direct refutation.** No probabilistic hedge protects it.

What would kill the program at this stage

K5.1 T1 FAILS: the model does not beat the baselines on held-out group order, with the surrogate calibrated.
CONSEQUENCE: the storage law does not describe group dynamics on this timescale with these proxies. Report and stop, or return to proxy selection with a new preregistration.

K5.2 A CLEAN COROLLARY 3.4 VIOLATION. A cluster with measured h_t well below zero grows its independently measured order proxy substantially, with eta, C, and N unchanged and the growth replicated in the held-out window.
CONSEQUENCE: the sustainability ceiling is false as stated. This is the single most consequential possible finding of the whole programme, because Corollary 3.4 is described in CANON as the most important sociological statement in the book. It would have to be withdrawn.

K5.3 T2 SUCCEEDS BUT T1 FAILS, REPEATEDLY.
CONSEQUENCE: the framework is a heuristic that generates useful interventions, not a theory that predicts. Real, respectable, and to be printed in exactly those words: a design manual, not a science.

K5.4 T3 FAILS ENTIRELY. No relation between internal and group parameters, in any cluster, on any measure.
CONSEQUENCE: Chapter 50's correspondence is an analogy of form, not of substance. CANON_INTERNAL §0's thesis reduces to a claim about MODEL FORM only, which is what CANON §1.4 already says it is. Survivable, and a large retraction of the volume's ambition.

K5.5 HARM. Any cluster in which the intervention measurably reduces members' ability to decline participation, exit, or dissent.
CONSEQUENCE: stop immediately. CANON §5, refusal 4 is not a methodological preference and does not trade off against a result.

Scale and timescale

clusters	24 (12 per arm)
members	20 per cluster; 480 total
periods	52 weekly
observations	24 x 20 x 52 = 24,960 member-periods
participant burden	4 minutes per member-period = 1,664 participant-hours
fieldwork	52 weeks
total stage	30 months: 6 site recruitment, 6 setup, 12 fieldwork, 6 analysis and dissemination

researcher and coordinator hours:

```
site coordination (24 x 52 x 2 h)              2496.0
site recruitment (24 x 40 h)                    960.0
instrument deployment and IT                    400.0
analysis                                        700.0
writing and dissemination                       300.0
------------------------------------------------------
TOTAL                                          4856.0
```

56.7 The kill register

Every kill condition in one table, so no stage is entered without its exit already written down.

Code	Stage	Kill condition	Consequence
K1.1	1	No shared unit for D and P	Moment equation demoted to heuristic ordering
K1.2	1	No instrument reaches session reliability 0.80	Programme cannot proceed
K1.3	1	Limiting cases have no recognisable referent	FVP-3 fails; formula has no phenomenon
K1.4	1	Y is a repackaging of U	Return to Stage 0 or stop
K2.1	2	No stable exponent across sessions	No scale; FVP-2 fails
K2.2	2	Cross-modality transitivity violated	Not a ratio scale; decisive
K2.3	2	No true zero for D or P	Division impermissible; demote in print
K2.4	2	D and P in incompatible units	As K1.1, quantified
K3.1	3	Any of the three rules fires in most subjects	Storage law does not describe internal dynamics
K3.2	3	Split-half reliability of b_hat below 0.50	eta is a window property; Stage 4 forbidden
K3.3	3	Systematic same-direction drift in every subject	Constant-parameter postulate fails
K3.4	3	Design controls failed; subjects inferred hypothesis	Data uninterpretable (52.7 per cent false-pass rate)

Code	Stage	Kill condition	Consequence
K4.1	4	Between-person spread does not clear the noise floor	delta and eta are not person parameters
K4.2	4	Rank order of subjects disagrees between labs	**FVP-7 fails. Programme stops.**
K4.3	4	Between-lab disagreement is a pattern, not a constant	As K4.2
K4.4	4	Suite passes in one lab, fails in the other	Do not proceed; find the difference or stop
K5.1	5	Model does not beat baselines on held-out group order	Storage law does not describe group dynamics here
K5.2	5	Clean Corollary 3.4 violation	Sustainability ceiling withdrawn
K5.3	5	Intervention works, model does not predict	Design manual, not a science; print those words
K5.4	5	No internal/external relation at all	Correspondence is form only, per CANON §1.4
K5.5	5	Ability to decline reduced in any cluster	Stop immediately; non-negotiable

Twenty-one ways to die: four terminal for the whole enterprise (K1.1 in its strong form, K3.2, K4.2, K4.3), and one — K5.5 — that stops the work regardless of what the data say.

56.8 Survival arithmetic

The following probabilities are **the author's declared priors, badged as priors and not as evidence.** They are stated so the expected cost can be computed at all; a reader who disagrees should substitute their own and redo the arithmetic, which takes one line.

```
  stage                       months   participants   P(fail here)   hours
  1  pilot operationalisation      7          1            0.10        376.5
  2  ratio-scale construction      4          6            0.10        778.9
  3  single-subject stability     12          3            0.15        831.5
  4  cross-subject recovery       18         80            0.35       1712.0
  5  field trials, group scope    30        480            0.20       4856.0
  ------------------------------------------------------------------------
  TOTAL (strictly sequential)     71        570                       8554.8

  SURVIVAL CURVE
```

```
after stage 1   P(still alive) = 0.9000
after stage 2   P(still alive) = 0.8100
after stage 3   P(still alive) = 0.6885
after stage 4   P(still alive) = 0.4475
after stage 5   P(still alive) = 0.3580

P(programme dies at or before stage 4) = 0.5525
```

Under these priors the program is more likely than not to be dead before the field trials begin, and its probability of reaching a completed Stage 5 is 35.8 per cent. That is not pessimism. It is what a research program looks like when its kill conditions are written honestly beforehand, and a program with a high stated completion probability has almost certainly written them too loosely.

The heaviest prior — 0.35 at Stage 4 — is deliberate. Stage 4 is where a formalism that describes meets the demand that it measure. CANON_INTERNAL §6 names FVP-7 the gate most often skipped; **it is most often skipped because it is most often failed.**

56.9 Resourcing

```
researcher and coordinator hours        8,554.8
at 1,600 productive hours per FTE-year      5.35 FTE-years
calendar, strictly sequential                 71 months = 5.92 years
calendar, with 30 per cent overlap of
  stages 2–3 and 4–5                          61 months = 5.08 years
implied continuous staffing        5.35 / 5.92 = 0.90 FTE
```

The program is one full-time researcher for six years. That is its actual shape: one person for six years, or two for three, with participants supplying a further 2,791 hours of their own time. Participant burden by stage:

```
stage 1     56.3 participant-hours       1 person
stage 2     38.9                         6 people
stage 3     72.0                         3 people
stage 4    960.0                        80 people
stage 5  1,664.0                       480 people
-------------------------------------------------
TOTAL    2,791.2 participant-hours     570 people
```

Monetary cost is not stated as fact because a defensible rate is not available to this manual. **As a declared assumption and not a citation**, at fully loaded hourly rates of 60, 90, and 120 in some currency unit, 8,554.8 researcher-hours cost 513,290, 769,935, and 1,026,580. Participant compensation, equipment, and site costs are additional and are not estimated, because doing so would require numbers this manual does not have. The figure that matters to a funder:

```
5.35 FTE-years / 0.3580 = 14.93 FTE-years per COMPLETED programme
```

About fifteen FTE-years per completed program, because most attempts stop early. The stopping is the point: nine or ten of those years are spent learning that something does not work, which is the ordinary and correct ratio in any field that tests things.

56.10 Tripwires for a degenerating program

A program degenerates when its response to disconfirmation is to add assumptions that protect the core rather than predictions that expose it. Each item below is a tripwire: if it happens, the program has stopped being a research program and someone must say so.

```
T1  A PARAMETER BECOMES TIME-VARYING AFTER IT FAILS TO BE CONSTANT.
    Legitimate: preregister a time-varying specification as a NEW
    prediction, on NEW data, with its own kill condition.
    Degenerate: refit existing data with delta_t and report the gain.

T2  N IS RE-AUDITED AFTER THE RESULT IS SEEN.
    N enters through ln N and through F, so a post hoc re-audit adjusts
    the ceiling to fit the data. Audit N before, freeze it, report it.

T3  THE PERIOD LENGTH IS CHANGED BECAUSE THE RESULT IMPROVES.
    One of the strongest levers on any autoregressive fit. Changing it
    after seeing results is a specification search wearing the clothes of
    a methodological refinement.

T4  A SUBJECT OR CLUSTER BECOMES "ATYPICAL" AFTER FAILING.
    Exclusion criteria are legitimate only if written in field A3 of the
    preregistration before recruitment.

T5  THE PROXY IS SWAPPED FOR A BETTER-BEHAVED ONE.
    The most common single rescue. Y is redefined, the fit improves, and
    nothing in the write-up records that a search occurred.

T6  A FAILED STAGE IS RECLASSIFIED AS A PILOT.
    A stage becomes a pilot when it was preregistered as a pilot, not when
    it produced an unwelcome number.

T7  THE KILL CONDITION IS RENEGOTIATED IN THE LIGHT OF THE DATA.
    The most serious of the seven, because it is the one that dissolves
    all the others. A kill condition altered after the data are seen was
    never a kill condition.
```

The countermeasure for all seven is structural, not moral: publish the kill conditions with the preregistration, before the data exist, where the investigators cannot edit them. A program that has done that can still degenerate, but not quietly.

56.11 What partial success looks like

Programs rarely die cleanly or succeed completely. Five partial outcomes, named in advance so they are recognised as outcomes rather than argued about as failures:

```
PS1  STORAGE LAW SURVIVES, MOMENT EQUATION DOES NOT.
     K1.1 or K2.3 fires; delta and eta recover on a directly measured Y.
     Parts VIII-IX stand; Part VII is demoted to diagnostics. The most
     probable partial outcome.
```

```
PS2  PARAMETERS ARE SPECIES-TYPICAL, NOT INDIVIDUAL.
     K4.1 alone fires. The framework describes a process everyone runs
     with similar constants. Useful, publishable, and requires rewriting
     every "your delta" in the volume.

PS3  IT WORKS INTERNALLY AND NOT EXTERNALLY, OR THE REVERSE.
     K5.4 fires. Two separate models with the same algebra, which is
     precisely what CANON §1.4 already claims and no more.

PS4  IT PREDICTS BUT DOES NOT GUIDE.
     T1 succeeds, T2 fails. A descriptive theory with no demonstrated
     intervention value. Honest and unglamorous.

PS5  IT GUIDES BUT DOES NOT PREDICT.
     T2 succeeds, T1 fails -- K5.3. A design manual whose heuristics work
     for reasons other than those stated. Print it in those words: a
     manual that works for unknown reasons is still a manual, and is not
     a science.
```

PS5 is the outcome most likely to be misreported, because it is the one where everything appears to be going well. Working interventions are the most persuasive and least diagnostic evidence available; Chapter 55's apparatus exists to keep those two apart.

56.12 What the program would not establish even if every stage succeeded

CANON §5's four refusals are not suspended by success.

```
NOT anything about physics, thermodynamics, or consciousness. Paper A is
pure mathematics and makes no empirical or metaphysical claim; Paper E
proves no negentropy law.

NOT that surplus is information in the Shannon sense. Paper D declines
that identification explicitly.

NOT the conversion of Volume Two's badged postulates into theorems. A
postulate that survives testing is a well-tested postulate.

NOT any intervention that removes a person's ability to decline. That
refusal is prior to the evidence and unaffected by it.
```

A program that cleared every gate would establish exactly this: that `delta` and `eta` are measurable person-level and group-level constants, in ratio-scaled units, reproducible across laboratories, whose values predict held-out order better than persistence, a linear input, autoregression, and the input's own timing scrambled. That is a great deal, and much less than the framework's ambitions. Saying so is what keeps the ambitions honest.

The Specification

S56.1 Stage 0's four preconditions (P1–P4) are in writing before any recruitment. Costless refusal is among them.

S56.2 Stages run in order. No stage begins before its predecessor's advance conditions are met and published.

S56.3 Every stage's kill conditions are published with its preregistration, before data exist, where the investigators cannot edit them.

S56.4 Stage 1 runs at least 45 sessions — 41 is the minimum at which an observed test-retest r = 0.70 has a 95 per cent lower bound above 0.50 — and reports session reliability with the item count required to reach it (Spearman-Brown).

S56.5 Stage 1 demonstrates or refutes a shared unit and a common zero for D and P, and prints whichever it found.

S56.6 Stage 2 spans at least 2.0 decades of stimulus range. Widening the range takes precedence over lengthening the session.

S56.7 Stage 2 collects enough judgments for SE(exponent) <= 0.02 (n >= 469 per exponent at a 2.0-decade span) before any transitivity claim is made.

S56.8 Stage 2's transitivity test states its detectable threshold explicitly alongside its result.

S56.9 Stage 2 demonstrates a true zero for D and for P, or the Moment equation is demoted in print per Chapter 53.

S56.10 Stage 3 samples daily, not twice weekly. The precision-per-week arithmetic is decisive and is reported.

S56.11 Stage 3 runs at least 240 periods per subject, split 120/120, and reports split-half reliability of a_hat and b_hat.

S56.12 Stage 3 reports the instrument's rating noise and the implied attenuation of delta and eta from Chapter 55 §55.9.

S56.13 Stage 3 states its power to detect within-person parameter change before claiming any such change.

S56.14 Stage 4 is run by two laboratories independently from the published protocol. A single-lab Stage 4 does not satisfy FVP-7 and must not be described as if it did.

S56.15 Stage 4 simulates the estimation noise floor at its own n and T and compares the observed between-person spread against it explicitly, before any individual-differences claim.

S56.16 Stage 4 publishes the reliability of the person-level estimate and the ceiling it places on correlations with external variables.

S56.17 Stage 4 classifies any between-lab disagreement as either a single scaling constant (calibration) or a pattern (theory failure), and prints which.

S56.18 Stage 5 runs both T1 (prediction) and T2 (intervention). Neither alone is sufficient.

S56.19 Stage 5 states its powered effect size given its cluster count and assumed ICC, and claims only what that supports.

S56.20 Stage 5's time series is at least 52 periods with a held-out window of at least 18, so the surrogate test has resolution at B = 999.

S56.21 Stage 5 tests Corollary 3.4 as a prohibition: clusters with negative measured headroom are predicted not to grow.

S56.22 All seven degeneration tripwires (§56.10) are monitored by someone with standing to halt the programme.

S56.23 Null results are published with titles that say so.

S56.24 The survival priors of §56.8 are restated at each stage boundary and updated in public.

The Failure Modes

F56.1 Stage Skipping

Signature. A study estimating `eta` for forty people that never built a ratio scale, or a field trial launched without any single-subject evidence. Stages 4 and 5 are the interesting ones. Stages 1 and 2 are apparatus.

Fix. S56.2. The estimate produced by a skipped-stage study is a number in unknown units, and unknown units cannot be compared to anything, which means the study cannot succeed even in principle.

F56.2 The Under-Powered Ratio Scale

Signature. A transitivity test reported as passed, with exponent standard errors near 0.10. Magnitude estimation is quick per judgment, so the total is underestimated.

Fix. S56.7 and S56.8. At `SE = 0.10` per exponent, only a discrepancy above 0.37 is detectable — the test cannot fail, which means it cannot pass either.

F56.3 Sampling Too Slowly

Signature. Twice-weekly experience sampling adopted to reduce burden. It genuinely does reduce burden.

Fix. S56.10. `T = 120` at daily sampling is 3.9 months; the same T twice-weekly is 13.8 months. The extra ten months are ten months of parameter drift, attrition, and life events, all of which attack the constant-parameter postulate that Stage 3 is trying to test.

F56.4 The Single-Lab FVP-7

Signature. "Replicated across two independent samples" where both samples were collected by the same group with the same hands. It is much easier and it feels like replication.

Fix. S56.14. The whole content of FVP-7 is that the procedure survives a different group's reading of the written protocol. Two samples from one lab test sampling variability, which was never the question.

F56.5 Individual Differences From Estimation Noise

Signature. A per-subject `eta` table with a spread of 0.03 to 0.04 at 120 periods, read as personality. The numbers differ, so the people appear to differ.

Fix. S56.15. The simulated floor is 0.036 at `T = 120` and 0.048 at `T = 60`.

F56.6 The Endpoint-Only Field Trial

Signature. A cluster-randomised trial with a baseline and a follow-up and nothing between. It is the standard design and it is cheaper.

Fix. S56.20. The framework's claim is about a *trajectory* governed by a recursion. A two-point design cannot distinguish the storage law from any other process that ends higher, and it certainly cannot run a surrogate test.

F56.7 Kill-Condition Erosion

Signature. The published kill condition and the one discussed in the group meeting have drifted apart; the discrepancy is nobody's fault and everyone knows about it. Sunk cost, plus the fact that every individual relaxation is small.

Fix. S56.3 and S56.22. Publish the conditions where they cannot be edited, and give someone standing to halt.

F56.8 Success Theatre at Stage 5

Signature. The intervention worked; the prediction test is described briefly in a supplement. PS5. A working intervention is the most persuasive and least diagnostic result available.

Fix. S56.18 and PS5's printing requirement. If T2 succeeded and T1 failed, the sentence "this is a design manual whose heuristics work for reasons other than those stated" must appear.

F56.9 The Rescued Corollary

Signature. A cluster with negative headroom grew, and the response is to re-measure `C`. Corollary 3.4 is described in CANON as the most important sociological statement in the book, and important statements attract rescues.

Fix. T2 and T5 of §56.10, plus S56.21. Freeze `C`, `eta`, and `N` before the window opens. If the violation survives a frozen measurement and replicates in held-out data, the corollary is withdrawn. That is what withdrawal means.

F56.10 Measuring People Who Cannot Refuse

Signature. Participation is enrolled at the organisational level and individual opt-out carries a cost, however small. Cluster designs recruit clusters.

Fix. S56.1 and K5.5. CANON §5, refusal 4. This one does not trade off against statistical power, and a design that cannot meet it is a design that does not run.

Worked Example

One laboratory, one principal investigator at 0.9 FTE, start date month 0. Every figure below is the arithmetic of §56.2 through §56.9 laid on a calendar; **the schedule is a plan, not a record — no month of it has been executed.**

The schedule

MONTH	ACTIVITY	GATE / OUTPUT
0–2	Stage 0: write P1–P4. Ethics. Consent infrastructure with costless refusal.	entry condition
0–4	Stage 1: instrument development, three candidates each for D, P, E.	160 h
3–7	Stage 1 fieldwork: 45 sessions, 3/week, 1 subject + 1 parallel replication subject.	56.3 subject-h
6–7	Stage 1 analysis. Publish reliability, shared-unit argument, limiting-case audit.	K1.1–K1.4 GATE 1
7–9	Stage 2: apparatus for 3 calibration modalities. Modulus selection.	240 h
8–10	Stage 2 fieldwork: 6 subjects, 2,916	38.9 subject-h

```
        judgments each, 6 sessions per term.
10–11   Stage 2 analysis: exponents, transitivity,       K2.1–K2.4
        true-zero test. Publish.                         GATE 2
-------------------------------------------------------------------
11–13   Stage 3: data pipeline, surrogate                320 h
        calibration simulations at b = 0.
12–20   Stage 3 fieldwork: 3 subjects, 240 daily         72 subject-h
        periods, 4 prompts/day (34.3 weeks).             2,880 prompts
20–23   Stage 3 analysis: per-subject fits, freeze,      K3.1–K3.4
        held-out, split-half. Publish.                   GATE 3
-------------------------------------------------------------------
23–26   Stage 4: publish protocol. Recruit lab 2.        150 h
        Recruit 40 + 40 subjects.
26–30   Stage 4 lab 1 fieldwork: 120 daily periods.      480 subject-h
29–33   Stage 4 lab 2 fieldwork, overlapping.            480 subject-h
33–41   Stage 4 analysis: noise floor at own n and T,    K4.1–K4.4
        recovery, between-lab reconciliation. Publish.   GATE 4
                                                         MAKE OR BREAK
-------------------------------------------------------------------
41–47   Stage 5: site recruitment, 24 clusters.          960 h
47–53   Stage 5 setup, instruments, randomisation,       400 h
        preregistration of T1, T2, T3.
53–65   Stage 5 fieldwork: 52 weekly periods.          1,664 participant-h
65–71   Stage 5 analysis and dissemination.              K5.1–K5.5
                                                         GATE 5
-------------------------------------------------------------------
TOTAL   71 months = 5.92 years, 8,554.8 researcher-hours,
        570 participants, 2,791.2 participant-hours
```

With the overlap shown at months 29-33 and 41-47, the calendar compresses to 61 months, 5.08 years.

The decision points, with their arithmetic

Gate 1, month 7. Does any operationalisation of `D` and any of `P` share a unit with a common zero? If no, the program does not stop — it *narrows*: the Moment equation is demoted in print, the Self-Storage Law is refitted with `S_t = F(U_t)` (input variant A, Chapter 55 §55.1) instead of `M_t`, and Stages 2 through 5 proceed on the narrower claim. **This is the one gate whose failure produces a smaller program rather than no program**, and recognising that in advance is what stops a Stage 1 failure from being fought instead of accepted.

Gate 2, month 11. Transitivity, detectable threshold 0.0739 in the exponent given `SE = 0.02` per exponent. A measured `|a_AC − a_AB ∗ a_BC|` above it ends the ratio-scale claim.

Gate 3, month 23. Three subjects, three verdicts. Advance requires the three-rule suite cleared in at least two of three, and split-half reliability of `b_hat` at or above 0.80 in at least two of three. **The simulated benchmark at 120 periods per half is 0.846; a real figure near 0.85 clears, near 0.40 does not, and the region between 0.50 and 0.80 is the genuinely awkward one — for which the preregistered rule is to extend to 360 periods in a fourth subject rather than**

to argue.

Gate 4, months 33-41. Three numbers decide it:

```
observed between-person sd of eta    must exceed 0.0355  (noise floor at T=120)
between-lab rank correlation of eta   must exceed  0       significantly,
                                                        and lie at or below the
                                                  published reliability ceiling
                                                     (0.942 at between-sd 0.10)
between-lab mean ratio of eta        must equal   1/g for an independently
                                                   measured operationalisation
                                                    difference g, or be absent
```

If the rank correlation is near zero, the program stops at month 41 having spent 3,698.9 researcher-hours — 2.31 FTE-years — and having established something real and publishable: that `eta` as operationalised is not a person-level measurement. **That is a worthwhile purchase, and it is the single most probable outcome under §56.8's priors at 0.35. Budget for it as the expected case, not the feared one.**

Gate 5, month 71. T1, T2, T3, and the Corollary 3.4 prohibition. As scheduled — 24 clusters, 20 members — the trial is powered for a standardised effect of 0.50 and under-powered for 0.30 at `ICC = 0.05`, where 36 clusters would be needed. The preregistration therefore claims a large effect or claims pilot status, not both.

The honest summary a funder should be given

> *We propose a five-stage, 71-month program, requiring approximately 5.35 FTE-years of researcher time and 570 participants contributing about 2,791 hours, to determine whether two parameters of a formal model of stored order — a retention loss `delta` and a conversion efficiency `eta` — can be measured on ratio scales, recovered as person-level constants, and reproduced across laboratories. Twenty-one preregistered kill conditions are published in advance; four are terminal. Under our declared priors, which are priors and not evidence, the program has a 35.8 per cent probability of reaching a completed final stage and a 55.3 per cent probability of stopping at or before the cross-laboratory recovery stage. Expected cost per completed program, accounting for early stopping, is about 14.9 FTE-years. We will publish at every stage, including and especially when the stage kills the program, and null results will be titled as null.*

That paragraph is the chapter. Concrete, sequenced, every number computed from a stated assumption, naming the year at which the whole thing most probably ends and committing in advance to publishing that ending.

A program that cannot be killed at any stage is not a research program. This one can be killed at twenty-one places, is most likely to die at month 41, and says so on the first page of its own proposal.

Part XI — The Old Notation

Chapter 57 . Correspondence Tables as Relational Graphs

Part XI is the part of this manual that a careful reader will approach with the most suspicion, and that suspicion is correct and should be kept. What follows is not a departure from the formalism. It is an application of the formalism to a body of historical technique, and the only things it is permitted to say are things the formalism can carry.

The author's position — recorded here as his thesis, with its badge on — is that the ancient occult and alchemical formulas are these same relationships expressed in a different notation. That is a claim about *notation*, and it is falsifiable in the only sense a claim about notation can be: either the old structures translate into the modern objects without residue, or they do not. This chapter performs the translation for correspondence systems. Chapter 58 performs it for the sigil. Chapter 59 performs it for the memory palace and the ritual calendar. Chapter 60 draws the balance and says what survives.

Nothing in Part XI asserts any supernatural claim, any physical effect of a symbol, or any result that is not derivable from the formalism already established in Chapters 1 through 56. That sentence is a standing constraint, it is repeated in every chapter of this part, and any sentence in this part that appears to contradict it is an error to be reported and corrected.

What *is* claimed here is narrower and more interesting. A correspondence table is a mathematical object of a completely ordinary kind: a finite relation over finite sets. The Tree of Life is a graph with ten vertices and twenty-two edges. The tarot is a structured set of seventy-eight elements with an explicit product decomposition. Gematria is a map from strings to integers — which is what a hash function is, and calling it one is a technical description, not a compliment and not an insult. These readings are descriptive. They involve no interpretation, no belief, and no metaphysics. They are what those objects *are*, and once you have said so you can compute with them.

The engineering payoff is specific. Paper B's validity conditions require a category assignment before the surplus functional may be applied to anything at all. You cannot compute `N_eff` until you have said what the N categories *are*. **A correspondence table is precisely a category assignment, written down, fixed, and made auditable — which is why the tradition that produced them was doing, in a different notation, the first step of the procedure this manual specifies.**

The Law

57.1 Statement of the reading rule

READING RULE (Correspondence Tables). Let a correspondence system be given: a finite set of labels L and one or more finite relations among them. Then:

```
    a single-column list of N labels          IS   a finite set of size N
   a two-column table over the same labels  IS  a binary relation R subset of L1 x L1
     an m-column table, one row per label       IS   (m-1) functions L1 -
> L2, ..., L1 -> Lm
   a diagram of nodes joined by paths       IS  an undirected graph G = (V, E)
   an assignment of labels to categories    IS  a partition, i.e. a block decomposition
   a letter-value scheme summed over a word  IS  a hash h : strings -> integers
```

There is no inference in that block. Each line is a restatement in modern vocabulary of what the historical artefact literally contains. **The reading rule is a translation, not a theory, and it cannot be wrong in the way a theory can be wrong — it can only be inapplicable, which happens exactly when the historical artefact is too vague to have a definite extension.** That case is common and is treated as failure mode F57.5.

The status table for this chapter:

Element	Status	Warrant
The Kircher-arrangement Tree of Life has 10 vertices, 22 edges	**PROVED** (enumeration)	§57.4, computed
Its density against `K10` is 22/45 = 0.4889	**PROVED** (arithmetic)	§57.4, computed
Its diameter is 3, mean geodesic 1.5778, cyclomatic number 13	**PROVED** (computation)	§57.4, computed
15 of its 22 edges cross pillar boundaries	**PROVED** (enumeration)	§57.5, computed
The tarot is `22 + 4 x 14 = 78`, a set with a product substructure	**PROVED** (arithmetic)	§57.6
Gematria is a function from strings to integers	**PROVED** (definition)	§57.8
Gematria destroys 10.47 bits on a 234,454-word list	**PROVED** (measurement)	§57.8, computed
The weekday sequence follows from `24 mod 7 = 3` on the Chaldean order	**PROVED** (arithmetic)	§57.7, computed
A correspondence table satisfies Paper B's category requirement	**PROVED** (definition)	§57.10
That these systems were *built as* mnemonic technology	**HISTORICAL CLAIM** (well-supported)	§57.9
That the old notation and this manual's notation describe the same relations	**POSTULATE** (author's thesis)	§57.1
That any correspondence carries efficacy, force, or physical consequence	**REFUSED**	CANON §5

The last two rows are the whole difficulty of this part of the manual, and they are separated on purpose. The thesis in the penultimate row is a claim about *isomorphism of structure*. The refused row is a claim about *action in the world*. **Structural correspondence is not evidence of mechanism, and this manual asserts no mechanism.** The same sentence appears in Chapter 29 about the square law and the infinite square well, and it means exactly the same thing here.

57.2 The honesty clauses

These are mandatory and are restated wherever the reading rule is used in Chapters 58 to 60.

H57-1 — Description, not endorsement. Saying that the Tree of Life is a graph is true in the same way that saying a subway map is a graph is true. It commits you to nothing about subways. Every metric computed in §57.4 is a metric of the *diagram as drawn*, and would be unchanged if the diagram had been drawn to organise plumbing fixtures.

H57-2 — The arrangement is a choice, and other arrangements exist. The ten-node, twenty-two-path figure used throughout this chapter is the arrangement commonly called the Kircher tree, reproduced in Athanasius Kircher's *Oedipus Aegyptiacus* (1652) and standard in later European occult literature. It is not the only historical arrangement; the Cordoveran and Lurianic schemes differ, and nineteenth- and twentieth-century revisions differ again. **All graph metrics reported here are metrics of the specific edge list printed in §57.4 and must be recomputed for any other arrangement.** An unstated arrangement produces unstated metrics, which is failure mode F57.4.

H57-3 — Gematria is a hash, and that is not a compliment. The description is exact and it is also deflationary. A hash function maps a large domain onto a small codomain. It therefore collides, massively and by construction. **Two words sharing a gematria value is the expected behaviour of a many-to-one map, not evidence about the words.** §57.8 measures the collision rate so that the deflation is quantitative rather than rhetorical.

H57-4 — Counting is not meaning. That the tarot has 78 elements, or that the Hebrew alphabet has 22 letters and the Kircher tree has 22 paths, are arithmetic facts. The nineteenth-century assignment of the 22 trumps to the 22 letters — associated with Éliphas Lévi, *Dogme et Rituel de la Haute Magie* (1854-56), and elaborated by later societies — is a *chosen bijection*, one of $22!$ available, made by identifiable people at an identifiable date. It is data about the history of the assignment and about nothing else.

H57-5 — Everything downstream inherits Papers A to E's clauses. When §57.5 computes a surplus number from a pillar decomposition, that number carries every honesty clause of the Surplus Law: the triangle inequality fails for geometric triples; subadditivity holds on u only; the functional as deployed is a **lower bound** on the ensemble's information; Paper A is pure mathematics and makes no empirical or metaphysical claim; no thermodynamic consequence follows.

H57-6 — The audit runs in both directions. If the translation reveals that an old table is a well-formed partition, say so. If it reveals that the table has duplicate rows, undefined membership, or an unjustified ordering, say that too, in the same voice. §57.11 does the second, and finds a 40.37% ceiling overstatement in a small worked case. A translation that only ever flatters the source is not a translation.

57.3 What Paper B requires before you may compute anything

Paper B's validity conditions are the gate this chapter exists to service. Before the surplus functional `F(u) = ln(1 + (N−1)u)` may be applied to any real ensemble, four things must hold, and they must hold *in print*, not in the analyst's head:

```
   V1  EXHAUSTIVE.   Every element of the domain falls in some category.
  V2 EXCLUSIVE.   No element falls in two, or the overlap is declared and handled.
```

```
 V3  COUNTED.     N is stated, and the N categories are genuinely distinct.
V4  ORDERED.    Where an ordinal is used (Chapter 29's essence index), the ordering is
                    justified, not arbitrary.
```

A correspondence table is a physical artefact whose entire job is to satisfy V1 through V4 for a domain, and to keep satisfying them across generations of users who cannot consult the author. That is a strong statement about what these systems were for, and §57.9 defends it on historical grounds. It is also the reason this chapter belongs in an engineering manual: the same four gates govern a hospital's incident-category scheme, a warehouse's fault taxonomy, and a research group's list of domains, and all three fail in the same ways.

The Engineering

57.4 The Tree of Life is a graph, and here are its numbers

The *Sefer Yetzirah*, the earliest text in the tradition to use the phrase, speaks of thirty-two paths — ten and twenty-two. The later diagrammatic tradition renders these as ten nodes joined by twenty-two edges. Take that diagram at face value and compute.

Vertices, in the conventional numbering:

```
1 Kether     2 Chokmah    3 Binah      4 Chesed     5 Geburah
6 Tiphareth  7 Netzach    8 Hod        9 Yesod     10 Malkuth
```

Edges, the Kircher arrangement, listed in full so the computation is reproducible:

```
1–2   1–3   1–6
2–3   2–4   2–6
3–5   3–6
4–5   4–6   4–7
5–6   5–8
6–7   6–8   6–9
7–8   7–9   7–10
8–9   8–10
9–10
```

Count: 3 + 3 + 2 + 3 + 2 + 3 + 3 + 2 + 1 = 22. No duplicates; verified by set comparison.

Computed metrics:

```
|V| = 10                          |E| = 22
sum of degrees = 44 = 2|E|        handshake lemma holds
K10 edges = 45                    density = 22/45 = 0.488889
cyclomatic number = |E|−|V|+1 = 13    (independent cycles)
connected: yes                    diameter = 3
mean shortest path = 1.577778    distance histogram: {1: 22, 2: 20, 3: 3}
```

Degree sequence:

Node	Name	Degree
6	Tiphareth	8
7	Netzach	5
8	Hod	5
2	Chokmah	4
3	Binah	4
4	Chesed	4
5	Geburah	4
9	Yesod	4
1	Kether	3
10	Malkuth	3

Sorted: `[8, 5, 5, 4, 4, 4, 4, 4, 3, 3]`.

Three observations, each of which is a fact about the diagram and about nothing else.

First, the degree distribution has a single dominant hub. Node 6 carries 8 of a possible 9 incident edges; it is adjacent to every node except Malkuth. The next-highest degree is 5. In the vocabulary of Chapter 21 and of Albert, Jeong & Barabási (*Nature*, 2000), this is a hub topology: robust to random edge loss, fragile to targeted removal of the hub. Delete node 6 and the diameter rises and the graph's mean path length degrades sharply; delete any leaf-adjacent node and little changes. That is a structural property of the drawing.

Second, the graph is far from a tree. A tree on 10 nodes has 9 edges. This has 22, hence 13 independent cycles. Alexander's "A City is Not a Tree" (1965) distinguishes the tree — where subsets nest strictly — from the semilattice, where they overlap. **The diagram called the Tree of Life is not, in Alexander's sense, a tree; it is a densely cyclic graph with an overlap structure, and that is why it can encode more relations than a hierarchy of the same size.** This is a pleasant terminological accident and it is also the substantive point: 13 independent cycles is a lot of overlap to carry in ten nodes.

Third, the endpoint-to-endpoint structure is rich. Distance from node 1 to node 10 is 3, and there are exactly 3 distinct geodesics of that length. The number of *simple* paths (no repeated vertex) from 1 to 10 is **573**, computed by exhaustive enumeration. A system whose users are taught to "travel the paths" has 573 distinct non-repeating routes between its two extremes. As a memory structure, that is the property being exploited: many routes to the same destination means many retrieval cues for the same content, which is exactly the design principle behind the method of loci treated in Chapter 59.

57.5 The three pillars are a block decomposition, and it is measurable

The same diagram carries a standard partition into three columns:

```
   Pillar of Severity   {3 Binah,  5 Geburah,  8 Hod}          size 3
   Pillar of Mercy      {2 Chokmah, 4 Chesed,  7 Netzach}      size 3
  Middle Pillar       {1 Kether, 6 Tiphareth, 9 Yesod, 10 Malkuth}   size 4
```

This is a partition: exhaustive, exclusive, three blocks summing to ten. It satisfies V1, V2 and V3 of §57.3 without amendment. **It is therefore a legitimate block decomposition in the sense**

Paper A Theorem 4.1 requires, and the cross-block and within-block edge counts can be measured rather than assumed.

Measured:

```
within-pillar edges:   Middle 3,  Mercy 2,  Severity 2       total  7
between-pillar edges:                                        total 15
cross-block fraction:  15/22 = 0.681818

cross-pair breakdown:  Mercy--Middle     6
                       Middle--Severity  6
                       Mercy--Severity   3

saturation:  Severity   2 of max 3 internal edges
             Mercy      2 of max 3
             Middle     3 of max 6
```

Two facts fall out. The blocks are internally *sparse* — no pillar is a complete subgraph — and the wiring is overwhelmingly *across* blocks, 68.2% of all edges. The two outer pillars connect to the middle six times each but to one another only three times. The diagram routes its cross-side traffic through the centre.

Now the translation into Paper A. Theorem 4.1, for a pure block vector, gives

```
u = u_cross + u_div       u_cross = beta^2       u_div = alpha^2 sin^2 gamma

 max(f(u_cross), f(u_div))  <=  f(u)  <=  f(u_cross) + f(u_div)
```

where `beta^2` is the between-category share and `alpha^2 sin^2 gamma` the within-category share. **Badge: the identification of the measured cross-block edge fraction with `beta^2` is a modelling choice, not a theorem.** It is the natural choice — the share of the structure that lives across categories — but Paper A does not tell you to make it, and a different operationalisation of `beta` gives different numbers. With `beta^2 = 15/22 = 0.681818` and `alpha^2 = 7/22 = 0.318182`, sweeping the within-block angle `gamma`, at `N = 10`:

gamma	u	N_eff	F (nats)	% of ceiling ln 10 = 2.302585
15 deg	0.703132	7.3282	1.991729	86.50%
30 deg	0.761364	7.8523	2.060803	89.50%
45 deg	0.840909	8.5682	2.148056	93.29%
60 deg	0.920455	9.2841	2.228302	96.77%
90 deg	1.000000	10.0000	2.302585	100.00%

And the subadditivity bracket, verified numerically at a representative point (`beta^2 = 0.68`, `alpha^2 = 0.32`, `gamma = 30 deg`, `N = 10`):

```
   u_cross = 0.6800   u_div = 0.0800   u = 0.7600   N_eff = 7.8400
max(f(u_c), f(u_d)) = 1.962908 <= f(u) = 2.059239 <= f(u_c)+f(u_d) = 2.505232   OK
```

The bracket holds at every sampled point, as it must. Note what the table shows: **once the cross-block share is above two thirds, the within-block angle barely matters — the figure is**

already at 86.5% of its ceiling at `gamma = 15 deg`. That is strict concavity doing its work, and it is the same lesson as Chapter 8: the first crossing purchases most of what crossings can buy, and the rest is glazing bills.

The honest coda: `ln 10 = 2.302585` nats is a *hard ceiling* on surplus per interaction for any ten-category scheme (Paper D, Corollary 2.2). No arrangement of paths raises it. Adding an edge to the diagram does not add a category. **The ceiling is set by `N`, and `N` is set by the table, not by the drawing.**

57.6 The tarot as a structured 78-element set

The deck is not a bag of 78 things. It has an explicit product structure, and the structure is the interesting part.

```
majors (trumps)          22
minors                   56  =  4 suits x 14 ranks
                             =  4 x (10 pips + 4 courts)
total                    78  =  22 + 40 + 16
```

Computed combinatorics:

```
naming one card:        log2(78)    = 6.285402 bits   ln(78) = 4.356709 nats
 unordered pairs C(78,2)                    = 3,003
 unordered triples C(78,3)                  = 76,076
ordered 3-card spread 78*77*76        = 456,456       log2 = 18.800116 bits
 ordered 10-card spread P(78,10)            = 4,566,176,969,818,464,000
```

The product structure is exact in the log register, which is the register everything in Volume Two operates in:

```
log2(4) + log2(14) = 2.000000 + 3.807355 = 5.807355 = log2(56)
```

That identity is trivial arithmetic and it is also the whole design principle of a factored code. **A suit-and-rank code lets a user reconstruct 56 identities from 4 + 14 = 18 memorised labels instead of 56, and the log-additivity is why it costs nothing in addressing power to do so.** Chapter 33 developed the same move for essence sets; here it appears in a fifteenth-century object.

Block structure over the minors, taking the four suits as categories:

```
all minor pairs    C(56,2) = 1,540
within-suit pairs  4*C(14,2) = 364
cross-suit pairs             = 1,176      cross fraction = 0.763636
```

76.4% of minor-arcana pairs cross a suit boundary. Compare 68.2% for the Tree's pillar decomposition. Both structures put most of their relational mass across categories rather than within them, which is the geometry Paper A rewards — `u_cross = beta^2` is the between-category term and it is the term that survives when within-category variation collapses.

One arithmetic tie-in to Chapter 29, offered as a structural remark and nothing more. The Essence Law's Result Two states that an ordered set indexed `1..N` yields exactly `ceil(N/2)` pitch classes, because pitch class depends only on the odd part of the index. Computed: for `N = 22`, 11 classes; for `N = 78`, 39 classes. **A 78-element ordered scheme has 39 distinct essence pitch classes,**

exactly half its size, and that is a hard structural limit of the indexing, not a property of the deck. It carries every clause of Chapter 29, including H1: essence, not resonance.

57.7 The sevens, and the week as an arithmetic consequence

The planet/metal/day table is the cleanest correspondence artefact in the European tradition because all three columns have exactly seven entries and the assignment is a bijection.

Day	Planet	Metal
Saturday	Saturn	lead
Sunday	Sun	gold
Monday	Moon	silver
Tuesday	Mars	iron
Wednesday	Mercury	quicksilver
Thursday	Jupiter	tin
Friday	Venus	copper

As a relational object: two independent bijections on 7 labels. The information content of pinning them down:

```
7! = 5,040 bijections      log2(5,040) = 12.299208 bits per column pairing
 a 3-column table pins 2 bijections:  2 * 12.299208 = 24.598416 bits
```

Twenty-four and a half bits is the entire informational content of a table that a tradition maintained for well over a millennium. That is a small number, and it is *supposed* to be a small number: the artefact's job is to be *stable and transmissible*, not to be dense. Compare the mnemonic-load computation in §57.9.

Now the part that is genuinely a derivation. The order of the days of the week looks arbitrary and is not. Take the seven classical planets in Chaldean order — descending apparent orbital period, which is the order they were ranked in by the sidereal periods known to antiquity:

```
Saturn, Jupiter, Mars, Sun, Venus, Mercury, Moon
```

Assign rulership of the hours of the day cyclically in that order, one planet per hour. The day is named for the ruler of its first hour. A day has 24 hours, and

```
24 mod 7 = 3
```

so the ruler of the first hour of the next day is **three steps on** in the Chaldean order. Stepping by 3 from Saturn and reading off:

```
Saturn -> Sun -> Moon -> Mars -> Mercury -> Jupiter -> Venus -> Saturn
```

which is

```
Saturday, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and back to Saturday
```

Computed and verified: the eighth step returns to Saturn, the cycle closes, and every planet is visited exactly once. The surviving English names carry the Germanic interpretatio for four of the seven (Tiw for Mars, Woden for Mercury, Thor for Jupiter, Frigg for Venus); the Romance languages

kept the Latin directly — *mardi*, *mercredi*, *jeudi*, *vendredi*. The earliest surviving explanation of the sequence by the planetary-hour argument is in Cassius Dio's *Roman History*, Book 37.

The weekday order is the star polygon {7/3}, and `gcd(7, 3) = 1` is exactly why the week is a single closed cycle that touches every planet rather than two or three disjoint sub-weeks. That is Chapter 58's theorem arriving a chapter early, in a piece of civil infrastructure that seven-day weeks are still keeping. If the day had been divided into a number of hours congruent to 0, 7, 14 or 21 modulo 7, the construction would have degenerated: every day would have had the same ruler. If it had been divided into 21 hours, `21 mod 7 = 0`, and the week would not exist. The week is a `gcd` result and the arithmetic is checkable in a minute.

57.8 Gematria is a hash function

Here is the operation. Assign each letter of an alphabet an integer. To evaluate a word, sum the integers of its letters. The standard Hebrew scheme (*mispar hechrachi*) assigns:

```
aleph 1   bet 2    gimel 3   dalet 4   he 5     vav 6    zayin 7   chet 8   tet 9
yod 10   kaf 20   lamed 30  mem 40    nun 50   samekh 60  ayin 70  pe 80   tsadi 90
 qof 100    resh 200 shin 300   tav 400
```

Twenty-two letters; units 1-9, tens 10-90, hundreds 100-400; the values sum to 1,495.

```
h : strings over an alphabet A  ->  non-negative integers
h(w) = sum over letters of w of value(letter)
```

That is a hash function. A hash function is a map from a large or unbounded domain into a smaller codomain, computed by a fixed deterministic rule. The definition contains no requirement that the map be interesting, only that it be a map. Gematria satisfies it exactly. Saying so is a literally accurate technical description of the operation, and it carries no claim whatever that the map is meaningful.

The description is useful because it makes the properties computable. A hash has three properties that engineers care about, and this one's scores are measurable.

Property 1 — order invariance. Gematria is a *multiset* hash: h is a sum, and addition commutes, so any permutation of a word has the same value. Anagrams collide by construction, and so does any rearrangement, meaningful word or not. Measured, purely from the permutation effect:

word length L	strings over 22 letters	distinct multisets	permutation collisions alone
2	484	253	1.913x
3	10,648	2,024	5.261x
4	234,256	12,650	18.518x
5	5,153,632	65,780	78.346x

At five letters, the ordering information — which is most of what a written word *is* — has already been thrown away at a rate of 78 to 1, before any value-arithmetic collisions are counted.

Property 2 — codomain size. The value range grows linearly in word length while the domain grows exponentially, so the collision rate grows without bound. Computed exactly over the Hebrew value set:

L	strings 22^L	distinct achievable sums	mean preimage	bits in	bits out
2	484	178	2.719	8.92	7.48
3	10,648	577	18.454	13.38	9.17
4	234,256	976	240.016	17.84	9.93
5	5,153,632	1,375	3,748.096	22.30	10.43
6	113,379,904	1,774	63,912.009	26.76	10.79

The input column climbs by 4.46 bits per added letter. The output column climbs by less than a bit. **By six letters the average value is shared by roughly sixty-four thousand strings.**

Property 3 — no avalanche. A cryptographic hash is designed so that changing one input bit changes about half the output bits. Gematria does the opposite: changing one letter changes the value by exactly the difference of the two letter values, a small, signed, entirely predictable amount. It is a *checksum*, and a weak one — it will catch a single-letter substitution but not a transposition, which is the most common scribal error there is.

Now the measurement that settles the matter, performed on a real word list rather than on a model. Using English simple gematria (A = 1 through Z = 26) over the macOS `/usr/share/dict/words` list (`web2`), restricted to purely alphabetic ASCII entries:

```
distinct words                    234,454
distinct hash values                  305
value range                         1..328
mean bucket size                  768.702
largest bucket                value 102 -> 2,353 words
top five buckets      102:2353  107:2325  97:2305  100:2295  115:2272
singleton buckets (no collision)        12
fraction of words colliding       0.999949

input entropy   log2(234,454)  =  17.839 bits
output entropy  H(hash)        =   7.371 bits
ceiling         log2(305)      =   8.253 bits
DESTROYED                      =  10.468 bits
```

Gematria on this list destroys 10.47 of the 17.84 bits that identify a word, and 99.995% of words share their value with at least one other word. The value 102 alone is shared by 2,353 words, beginning `abrogation, absurdly, acanthous, acatamathesia, acclimatize, achromatin, acontius, acrophonic, acroscopic, acroterial, actability, actionably, actionize, acupress`.

The engineering consequence is exact and it is the point of the section. **A shared gematria value is the default outcome, not a signal.** Discovering that two words share a value is discovering that a many-to-one map behaved as a many-to-one map. Under a uniform model, two arbitrary

words from a 305-value codomain collide with probability on the order of 1/305, and words are not uniformly distributed, which makes it worse: the modal buckets hold over two thousand entries each. Any argument of the form *these two terms share a value, therefore they are related* is a statement about the map's compression ratio and contains no information about the terms. That is failure mode F57.2, and it is the most common analytical error in the entire literature this chapter is reading.

The deflation cuts one way only. It says the map carries no evidential weight. It does not say the map is useless: a digest that reduces 17.8 bits to 7.4 is a genuine index, and an index is a real mnemonic aid for exactly the reason §57.9 gives. Cross-referencing by digest is how you find candidate material in a corpus you cannot search. What you may not do is treat the retrieval cue as a result.

57.9 What these systems were actually for

The claim of this section is historical and it is well supported, but it is a historical claim and is badged as one: **these correspondence systems were mnemonic and combinatorial technologies — tools for holding a large number of relations in a mind, in cultures where written notation was scarce, expensive, restricted, or dangerous to possess.**

The load calculation makes the case quantitative. To hold `n` items as an unstructured web of pairwise relations costs `n(n−1)/2` facts. To hold them as a single ordered cycle costs `n` successor facts. To hold them as a table with `m` columns costs `n` rows.

n	unstructured pairs `n(n−1)/2`	ordered cycle `n`	ratio
7	21	7	3.00x
10	45	10	4.50x
12	66	12	5.50x
22	231	22	10.50x
78	3,003	78	38.50x

At `n = 78` the structured encoding is 38.5 times cheaper to hold than the relation set it addresses, and the saving grows linearly in n. That is the whole reason to build a correspondence table, and it is a reason that survives complete disbelief in every metaphysical claim ever attached to one.

Four design features of these systems are recognisable as mnemonic engineering, and each has a modern counterpart already treated in this manual:

(a) Small, fixed, memorable cardinalities. Seven planets, seven metals, seven days. Ten sephiroth. Twenty-two letters. Four elements. Twelve signs. These sit in the range a person can hold and rehearse. Dunbar's number is contested and should be stated as contested, but no one contests that seven is easier than seventy.

(b) Factored codes. Suit-and-rank: 18 labels address 56 identities. Planet-and-hour: 7 labels address 168 hour-slots per week. The log-additivity shown in §57.6 is why factoring is free in addressing power and enormous in memory cost.

(c) Redundant retrieval paths. 573 simple paths from node 1 to node 10 in the Kircher tree. Many cues, one content. This is the same principle as the method of loci (Yates, *The Art of Memory*,

1966), which Chapter 59 treats as the storage law with space as the retention medium — high `eta`, low `delta`, because place is durable.

(d) Cross-modal binding. A correspondence table binds a name to a metal, a colour, a day, a scent, a number. Each additional column is another sensory handle on the same slot. This is cheap to add — one more row-entry per label — and it multiplies the retrieval cues.

(e) Combinatorial generativity. A table with `m` columns over `n` labels supports `n^2` ordered pairings, `C(n,2)` unordered ones, and arbitrary compositions across columns. The user does not need to have memorised the pairing; the pairing is *computed* from the table when needed. Lull's *Ars Magna* in the thirteenth century made this mechanical and explicit, with rotating concentric wheels that generate combinations of terms — a physical combinatorial engine.

None of (a) through (e) requires anything to be true about planets. They are the properties of a well-designed index in a world without indexes. **The systems are engineering artefacts of a severe memory constraint, and they are good ones; when the constraint was lifted by cheap print and then by search, the artefacts stopped being infrastructure and became objects of study, which is exactly what Star & Ruhleder (1996) predict — infrastructure becomes visible on breakdown.**

57.10 The join: a correspondence table is a category assignment

Now the connection this chapter exists to make.

Paper A's decomposition (Theorem 4.1) is stated for a *block* structure: the ambient space splits into blocks `H_1, ..., H_m`, a vector's mass divides between blocks, and the surplus separates into a cross-block term and a within-block term. Paper B's validity conditions (§57.3) say you may not apply the functional until the blocks are declared. Chapter 29 says the essence index is meaningless without a justified ordering of categories.

Every one of those requirements is a requirement for a **table**: rows are categories, and the assignment of things to rows is the block structure.

```
correspondence table            Paper A / Paper B object
-------------------------------------------------------------
the set of row labels           the category set, cardinality N
N                               the surplus ceiling ln N
assignment of items to rows     the block decomposition
within-row variation            u_div = alpha^2 sin^2 gamma
between-row relations           u_cross = beta^2
the row ordering (if used)      the essence index n (Chapter 29)
a column                       a function from categories to attributes
an unassigned item            a V1 violation: the table is not exhaustive
an item in two rows            a V2 violation: declare it or repartition
two rows naming the same thing  a V3 violation: N is inflated (see §57.11)
```

This is the substantive content of the author's thesis in the case of correspondence systems: the old artefact and the new formalism are doing the same job at the same point in the procedure, and the translation costs nothing on either side. It is not an argument that the old artefact is correct. A table can satisfy V1 through V4 perfectly and still assign the wrong attributes; being well-formed is not being true. It is an argument that the old artefact is the *right kind of object*

— that when a fifteenth-century practitioner drew up a table of seven rows and four columns, the thing produced was a category scheme with a declared cardinality, which is precisely what a modern analyst must produce before computing `N_eff`.

The practical use is immediate. If you inherit a category scheme — a fault taxonomy, a diagnostic code list, a set of "customer segments" — you can run it through V1 to V4 exactly as you would run a historical table, and the audit will find the same faults.

57.11 Auditing N: the ceiling inflation fault

The most valuable thing the translation buys is a number that the tradition never computed: whether the table's `N` is real.

Take a seven-row table. Its declared ceiling is `ln 7 = 1.945910` nats, `2.807355` bits. Now suppose the audit finds that three of the seven rows are not distinct — that they name the same underlying category under different labels, which is the single most common defect in inherited taxonomies. The real cardinality is 4.

```
apparent ceiling   ln 7 = 1.945910 nats
real ceiling       ln 4 = 1.386294 nats
overstatement      0.559616 nats = 40.37% of the real ceiling
```

A scheme with three duplicate rows out of seven overstates its own surplus ceiling by 40.4%, and every downstream number — `N_eff`, `F`, headroom `h`, the storage law's growth condition — inherits that error at full size. This is the same fault as Chapter 11's pseudo-diversity, arriving from the taxonomy side rather than the roster side, and it has the same signature: an organisation that believes it holds seven kinds of thing and behaves accordingly, while its actual variety is four.

Ashby's Law of Requisite Variety is the reason this matters and not merely a bookkeeping irritation. Only variety absorbs variety. A control scheme with a genuine `N` of 4 cannot regulate a disturbance space of 7 kinds, and the table that says it has 7 rows will conceal the deficit until the fourth unhandled disturbance arrives.

The diagnostic is mechanical: for every pair of rows, ask whether any observable item has ever been assigned to one and not the other, and whether any decision differs between them. If neither, the rows are duplicates. Merge them and recompute the ceiling before doing anything else.

The Specification

Numbered, checkable requirements for building or auditing any correspondence table — historical, inherited, or new.

S57.1 — Publish the extension. The table must list its row labels exhaustively, in a fixed order, in one place. A correspondence system whose row set is "and so on" is not a category scheme and no quantity may be computed from it. Reject.

S57.2 — Declare N and derive the ceiling. State `N` and print `ln N` and `log2 N` beside it. Every downstream quantity is bounded by `ln N` (Paper D, Corollary 2.2) and the ceiling should be visible at the top of the artefact, not buried.

S57.3 — Pass V1 through V4 in print. Exhaustive, exclusive, counted, ordered-if-ordered. Each gate gets a line and a verdict. A failed gate demotes the artefact from *scheme* to *checklist* and the formalism may not be applied to it.

S57.4 — Run the duplicate audit before the first computation. For all `C(N,2)` row pairs, apply the §57.11 test. Report the number of merges performed and the corrected `N`. Print the overstatement in nats if any merge occurred.

S57.5 — Provide a residual row. Every scheme over an open domain requires an explicit `OTHER / NOT YET CATEGORISED` row, and the count in that row must be reported every period. A scheme with no residual row is not exhaustive; it is merely silent about what it drops. The rate of growth of the residual is the scheme's obsolescence signal.

S57.6 — State the arrangement whenever you state a metric. Any graph number — degree, diameter, density, cross-block fraction — is a number about one edge list. Print the edge list or cite it precisely. See H57-2.

S57.7 — Never treat a digest as a result. Any letter-sum, checksum, numerological value, or other many-to-one map may be used as a *retrieval index* and may not be used as *evidence of relation*. Where such a value is reported, report the collision rate of the map alongside it. See §57.8 and F57.2.

S57.8 — Justify the ordering separately from the membership. V3 and V4 are independent gates. A defensible set of categories with an arbitrary ordering yields arbitrary essence indices (Chapter 29, H4). If no ordering is defensible, use the scheme unordered and do not compute an index.

S57.9 — Measure the cross-block fraction; do not assume it. Count the relations that cross a category boundary and the relations that stay inside one, and report the fraction. This is the empirical handle on `beta^2`, and treating the measured fraction as `beta^2` is a modelling choice that must be badged each time it is made.

S57.10 — Prefer factored codes. Where a scheme addresses `n = p*q` identities, encode it as `p + q` labels with two columns rather than `n` labels with one. The addressing power is identical by log-additivity; the memory and maintenance cost falls from `n` to `p+q`.

S57.11 — Build redundant retrieval paths deliberately. More than one route to each entry. Count them: the number of distinct simple paths between the scheme's extreme entries is a computable robustness figure, as it was for the Kircher tree (573).

S57.12 — Re-audit on a fixed cadence. Category schemes drift: the world adds kinds and the table does not. Set a period, put the audit on it, and record `N` at each audit. A `delta` on the table itself is real, and Chapter 59's periodic-forcing argument applies to maintaining it.

S57.13 — Keep the badge on. Where the artefact carries traditional attributions, say what is computed and what is inherited, in the same paragraph. A metric on a diagram is not an endorsement of the diagram's cosmology, and a reader must never have to guess which register a sentence is in.

The Failure Modes

F57.1 — PHANTOM COLUMN. *Signature:* the table has a column whose entries cannot be independently checked against anything outside the table. Adding a row to the table is the only way anyone ever learns a new entry in that column. *Mechanism:* the column is definitional rather than empirical; it carries zero bits about the world because it is a relabelling of the row. *Fix:* mark such columns as *nomenclature* and exclude them from any information or surplus computation. They may be retained as mnemonic handles under S57.10; they may not be counted as attributes.

F57.2 — HASH MYSTICISM. *Signature:* an argument of the form "these two terms share a numerical value, therefore they are related." Escalated form: a search over many pairs until a coincidence is found, reported without the search. *Mechanism:* a many-to-one map colliding as many-to-one maps do. Measured on a real word list, the English simple-gematria map has 305 output values for 234,454 words, mean bucket 768.7, and 99.995% of words collide. The finding is a property of the map's compression ratio. *Fix:* report the collision rate with every value. Apply S57.7 without exception. If an argument survives the disclosure that its coincidence had a one-in-a-few-hundred prior, it was not resting on the coincidence and should be restated without it.

F57.3 — CEILING INFLATION BY DUPLICATE ROWS. *Signature:* the declared `N` exceeds the number of rows that any decision actually distinguishes. No item has ever been assigned to row `i` rather than row `j`, and no action differs between them. *Mechanism:* `N` enters the formalism as `ln N`, the hard ceiling. Duplicates raise the nominal ceiling without raising real variety. Computed instance: 3 duplicates in 7 rows overstates the ceiling by 0.559616 nats, 40.37%. *Fix:* S57.4, before any other computation. Then re-derive headroom `h = eta ln N − delta Q − C` with the corrected `N`, and expect it to fall.

F57.4 — UNSTATED ARRANGEMENT. *Signature:* a graph metric is quoted — "twenty-two paths", "the diameter is three" — with no edge list, and the source diagram turns out to differ between editions. *Mechanism:* the metric is a function of the edge list. Different traditions draw different edges; the numbers are not comparable across them. *Fix:* H57-2 and S57.6. Print the edge list. Where two arrangements are in play, compute both and report the difference; the difference is itself informative about what the tradition treated as optional.

F57.5 — VAGUE EXTENSION. *Signature:* the scheme cannot answer, for a specific concrete item, which row it belongs to — and different competent users answer differently. *Mechanism:* V1 and V2 fail simultaneously. There is no partition, therefore no block decomposition, therefore no `u_cross` and no `u_div`. *Fix:* either sharpen the row definitions until inter-rater assignment is reproducible (this is FVP-7's cross-individual gate applied to a taxonomy), or demote the artefact in print to a heuristic list and stop computing with it. Both are acceptable outcomes; silence is not.

F57.6 — EDGE INVENTION. *Signature:* a relation is asserted between two labels because the diagram has room for it, or because symmetry would be prettier, rather than because a source or an observation carries it. *Mechanism:* the graph is being fitted to an aesthetic prior. This is over-framing, `D/P > 1`, at the level of the notation itself (Chapter 35). *Fix:* every edge carries a provenance token — source, observation, or *conjecture*. Compute all metrics twice, once with conjectural edges and once without, and print both. The gap is the scheme's exposure to its own aesthetics.

F57.7 — UNIVERSAL CORRESPONDENCE. *Signature:* the scheme claims to categorise every-

thing. Any new object can be placed in it, and no observation could ever fall outside it. *Mechanism:* a scheme that cannot fail to apply carries no information about what it applies to. This is FVP-6's falsification gate failing at the taxonomy level, and it is Chapter 37's test for a legitimate postulate. *Fix:* state, in print, a concrete item the scheme does *not* categorise, or an observation that would force a new row. If no such item can be named, the scheme is unfalsifiable and must be demoted to a language rather than a measurement instrument.

F57.8 — TREE-FLATTENING. *Signature:* a densely cyclic scheme is redrawn as a strict hierarchy "for clarity", and the overlaps are deleted. *Mechanism:* Alexander, "A City is Not a Tree" (1965). A semilattice flattened to a tree loses exactly the overlapping memberships that carried the cross-block relations. Concretely: the Kircher arrangement has 13 independent cycles; the spanning tree has 0. Flattening deletes 13 relations and the cross-block fraction collapses. *Fix:* refuse the flattening, or perform it only as a *view* over a preserved cyclic master, and label the view as lossy. Count the deleted cycles and report the count.

F57.9 — INHERITED N. *Signature:* the scheme's cardinality is a traditional number — seven, ten, twelve, twenty-two — and no one can say what would have happened if the domain had contained eight kinds. *Mechanism:* N was set by the notation rather than by the domain. The scheme then forces the domain into its cardinality, merging real distinctions or splitting artificial ones. *Fix:* derive N from the domain first, independently, then compare it to the traditional number and print both. Where they differ, the traditional number is a hypothesis about the domain and should be tested, not assumed. Where they agree, say by how much the agreement was constrained.

Worked Example

The maintenance department's fault taxonomy

A building-services department of fourteen staff maintains a research campus. It has an inherited fault taxonomy of **seven categories**, established a decade ago, used to route every incoming work order. The head of department wants to know whether the taxonomy is real, what it costs, and whether the department's stored capability is growing or decaying. We run the chapter end to end.

Step 1 — Publish the extension (S57.1). The seven rows, as inherited:

```
    1  ELECTRICAL           5  FABRIC / BUILDING ENVELOPE
    2  MECHANICAL / HVAC    6  PLANT FAILURE
    3  PLUMBING             7  EQUIPMENT FAULT
    4  CONTROLS / BMS
```

Step 2 — Declare N and the ceiling (S57.2).

```
    N = 7      ln 7 = 1.945910 nats      log2 7 = 2.807355 bits
```

Step 3 — Validity gates (S57.3).

```
   V1 EXHAUSTIVE  FAIL. No residual row. 11% of work orders in the last year were
                  filed under row 7 with free-text that matched none of rows 1–
6.
    V2 EXCLUSIVE    FAIL.  Rows 6 and 7 overlap and rows 2 and 6 overlap.
```

```
 V3 COUNTED      PENDING duplicate audit.
V4 ORDERED    N/A.  The taxonomy is unordered; no essence index is computed. Correct.
```

Step 4 — Duplicate audit (S57.4). For all `C(7,2) = 21` row pairs, ask: has any item ever been assigned to one and not the other, and does any decision differ?

```
(2, 6)  PLANT FAILURE vs MECHANICAL/HVAC  -- no routing difference; same two engineers;
                                          same parts store.             MERGE
(6, 7)  PLANT FAILURE vs EQUIPMENT FAULT  -- distinction is who owns the asset,
                                         not what the fault is.         MERGE
(5, 7)  FABRIC vs EQUIPMENT FAULT        -- distinct: different trade, different SLA.  kee
 all others                                                              keep
```

Two merges. Rows 2, 6 and 7 collapse into a single category. The corrected count:

```
corrected N = 4 + 1 = 5   ... but row 7's residual traffic must go somewhere.
```

Adding the mandatory residual row (S57.5) restores a fifth and sixth:

```
1  ELECTRICAL
2  MECHANICAL / PLANT / EQUIPMENT      (merged from old 2, 6, 7)
3  PLUMBING
4  CONTROLS / BMS
5  FABRIC / ENVELOPE
6  UNCATEGORISED  -- residual, counted and reported every period

corrected N = 6      ln 6 = 1.791759 nats
```

The overstatement carried by the inherited taxonomy, in the pure form computed in §57.11 where a seven-row scheme is really four:

```
ln 7 - ln 4 = 1.945910 - 1.386294 = 0.559616 nats = 40.37% overstatement
```

In this department's actual case, seven collapsing to six is milder:

```
ln 7 - ln 6 = 1.945910 - 1.791759 = 0.154151 nats = 8.60% overstatement
```

The department believed it had 7 kinds of fault and had 6, one of which is an admission of ignorance it had not previously been counting. The 8.60% is the honest correction and it is smaller than the textbook case because two of the three suspect rows genuinely merged into one another rather than into three separate existing rows.

Step 5 — Cross-block measurement (S57.9). Over the last 400 work orders, count how many required a person from more than one category to close:

```
single-category closures     262
cross-category closures      138
cross fraction               138/400 = 0.345
```

Badged modelling choice: take `beta^2 = 0.345`, hence `alpha^2 = 0.655`. This is the same choice made in §57.5 and it carries the same badge — Paper A does not require it.

Step 6 — Surplus. With `N = 6` and a within-category framing angle estimated at `gamma = 40 deg` (from the spread of independent diagnoses recorded on cross-category jobs):

```
u_cross = 0.345000
u_div   = 0.655 * sin^2(40 deg) = 0.655 * 0.413176 = 0.270630
u       = 0.615630
N_eff   = 1 + (6-1)*0.615630 = 4.078151
F       = ln(4.078151) = 1.405644 nats        ceiling ln 6 = 1.791759
utilisation = 1.405644 / 1.791759 = 78.45% of ceiling
```

Subadditivity check, as required:

```
f(u_cross) = ln(1 + 5*0.345000)   = ln(2.725000) = 1.002468
f(u_div)   = ln(1 + 5*0.270630)   = ln(2.353150) = 0.855756
max = 1.002468  <=  f(u) = 1.405644  <=  sum = 1.858224       OK
```

Interpretation, stated as the formalism permits: a six-category department whose staff genuinely hold non-parallel framings is running at 78.45% of the surplus its own taxonomy allows. The remaining 21.55% is bought only by raising u — more genuinely different framings, or more cross-category work — and it is capped at ln 6 no matter what. To go beyond ln 6 the department must legitimately add a category, which means finding a kind of fault it currently cannot name. That is what the residual row is for, and it is why S57.5 requires it.

Step 7 — Feed the storage law. The department's parameters, estimated from its own records:

```
    eta   = 0.42     fraction of interaction that becomes durable capability
                                    (documented procedures, trained second-
persons, updated asset records)
    delta = 0.11     annual decay: staff turnover 14%, plus procedure drift
      Q_t    = 5.6          current stored capability, in the same nat-
scaled units as eta*S
    C_t   = 0.31     coordination overhead per period
```

Headroom, the master gauge:

```
h_t = eta * ln N - delta * Q_t - C_t
    = 0.42 * 1.791759 - 0.11 * 5.6 - 0.31
    = 0.752539 - 0.616000 - 0.310000
    = -0.173461
```

Negative. h_t < 0 means that in this model no interaction geometry saves it (Corollary 3.4). The department cannot grow its stored capability by rearranging who works with whom, because even at the absolute ceiling ln 6, conversion times ceiling does not cover decay plus cost.

Note what the audit did to this verdict. With the *inherited* N = 7:

```
h_t = 0.42 * 1.945910 - 0.616000 - 0.310000 = 0.817282 - 0.926000 = -
0.108718
```

Still negative — but 37% less negative, and a manager reading the uncorrected figure would have been closer to believing that a modest push on u could close the gap. It could not, in either case, but the corrected figure shows the true size of the shortfall. **The audit did not change the sign of the verdict; it changed the size of the intervention the verdict requires, which is the number the department actually needs.**

The three admissible responses, in the order the formalism ranks them:

```
   raise eta   0.42 -> 0.52 :  h = 0.931715 - 0.926000 = +0.005715  marginal, positive
     cut  C          0.31 -> 0.18 :   h = 0.752539 - 0.616000 - 0.180000 = -
0.043461   still negative
        cut  delta        0.11  ->  0.07  :    h  =  0.752539  -  0.392000  -
0.310000 = +0.050539   positive
        cut  delta  AND  raise  eta      :    h  =  0.931715  -  0.392000  -
0.310000 = +0.229715   comfortable
```

The dominant lever is `delta`, which for this department means turnover and procedure drift — and that is the honest engineering conclusion. Reorganising the taxonomy was necessary to get an accurate number and was never going to be sufficient to fix the problem. **A correspondence table tells you what `N` is; it does not tell you what `delta` is, and `delta` was the binding constraint.**

What the worked example demonstrates

Every step above was performed on an ordinary building-services taxonomy, and every step is identical in form to what §§57.4 through 57.8 performed on a seventeenth-century diagram and a scheme of letter values. The gates are the same, the block decomposition is the same, the ceiling is the same `ln N`, and the storage law that consumes the result is the same equation.

That identity of procedure is the entirety of what this chapter claims. It is the author's thesis in its narrow, defensible form: **the old correspondence systems and this manual's category requirement are the same object at the same point in the same procedure, expressed in different notation.** The old notation carried it in seven rows and four columns, memorised, because there was no other way to carry it. The new notation carries it in a partition, a cardinality, and a cross-block fraction, because now there is.

Chapter 58 takes the next object in the old notation — the figure drawn on the page rather than the table held in the mind — and finds that it has an exact arithmetic invariant that the tradition drew for centuries without naming.

Chapter 58 . The Sigil as Nonlinear Map

This chapter contains the strongest formal result in Part XI, and it is genuine number theory. It is also the chapter where the author's framing of a sigil as a *computing map* is stated, and that framing is a **METAPHOR**. The author says so himself, this manual keeps the badge on throughout, and every quantitative statement below is separated from it by a line the reader can check.

The situation is unusual and worth naming before we begin. Most of the material in Part XI is translation: an old artefact turns out to be an ordinary mathematical object once you look at what it literally contains. This chapter is different. The star polygon has an **exact arithmetic invariant** that the tradition drew for two thousand years and never named, and once the invariant is stated it settles, in one line, a question the tradition handled by trial and error: which figures can be drawn without lifting the pen. The answer is a greatest common divisor. It is true, it is checkable in a minute, and it was always there in the figures.

A star polygon {N/k}, drawn by stepping k positions around N equally spaced points, decomposes into exactly gcd(N,k) separate closed paths, each of length N/gcd(N,k).

The hexagram is two triangles because `gcd(6,2) = 2`. The heptagram is one continuous line because `gcd(7,2) = 1`. The twelve-point figure of step three is three squares because `gcd(12,3) = 3`. None of this is interpretation; all of it is arithmetic, and §58.5 prints the complete table for `N = 3` through `N = 24` with every valid `k`.

Nothing in this chapter asserts any supernatural claim, any physical effect of a symbol, or any result not derivable from the formalism. That standing constraint from Chapter 57 governs here without amendment.

The Law

58.1 The GCD Theorem

THEOREM 58.1 (Decomposition of the star polygon). Let `N >= 3` and `1 <= k <= N−1`. Place `N` labelled points at the vertices of a regular N-gon, identified with the residues `Z_N = {0, 1, ..., N−1}`. Draw the figure `{N/k}` by starting at some point and repeatedly joining the current point `x` to the point `x + k (mod N)`, continuing until the start is reached, then beginning again at any unused point, until every point has been used. Let `g = gcd(N, k)`. Then:

```
(i)   the figure consists of exactly  g  closed paths (strokes)
(ii)  each stroke visits exactly      N/g  distinct points
(iii) each stroke has exactly         N/g  edges
(iv)  the whole figure has exactly    N    edges, for every k
(v)   the figure is a single closed path (UNICURSAL) if and only if  g = 1
```

Proof. The step map is `s(x) = x + k (mod N)`. Its orbits are the cosets in `Z_N` of the cyclic subgroup generated by `k`. In the additive group `Z_N`, the subgroup generated by `k` is

```
<k> = { 0, k, 2k, ... }  =  <gcd(N,k)>  =  { 0, g, 2g, ..., N−g }
```

— the standard fact that the set of integer combinations of `N` and `k` is the set of multiples of their greatest common divisor. That subgroup has order `N/g`. By Lagrange's theorem the number of

cosets is |Z_N| / |<k>| = N / (N/g) = g. Each coset is one orbit of s, hence one closed path, and each has N/g elements, giving (i) and (ii). Every point has exactly one outgoing and one incoming edge, so each stroke of N/g points carries N/g edges, giving (iii); summing over g strokes gives g * (N/g) = N edges, giving (iv). And (v) is (i) with g = 1. **QED**

Equivalently, and this is the formulation to hold in mind: the figure is the Cayley graph of Z_N with connection set {+k, −k}, and its connected components are the cosets of <k>. Every vertex has degree 2, so the graph is 2-regular, so every component is a cycle. **A 2-regular graph is connected if and only if it is a single cycle, which is why unicursality, Eulerianity and Hamiltonicity coincide for star polygons and all three reduce to gcd(N,k) = 1.**

Verified empirically before being trusted: the figure was traced by explicit walk for every N = 3..40 and every k = 1..N−1, a total of 779 cases, and the stroke count, stroke length and edge total agreed with the theorem in **every case, zero mismatches.**

The Schläfli notation {N/k} and the classification of gcd > 1 cases as *star figures* or *compounds* rather than star polygons proper is standard; Coxeter's *Regular Polytopes* is the usual reference.

58.2 The corollaries that do the work

COROLLARY 58.2 (Unicursality). {N/k} can be drawn without lifting the pen if and only if gcd(N,k) = 1. The number of such k in 1..N−1 is phi(N), Euler's totient. Since {N/k} and {N/(N−k)} are the same figure traversed in opposite directions, the number of *distinct* unicursal figures on N points with constant step is phi(N)/2 for N > 2.

COROLLARY 58.3 (Pen lifts). Drawing the complete figure requires exactly gcd(N,k) − 1 pen lifts. Unicursal figures require zero.

COROLLARY 58.4 (Compound reduction). For g = gcd(N,k) > 1,

```
{N/k}  =  g  disjoint copies of  {(N/g)/(k/g)}
```

and the inner figure is always unicursal, because gcd(N/g, k/g) = 1 by construction. Verified:

```
{6/2}  = 2 x {3/1}    two triangles
{6/3}  = 3 x {2/1}    three digons (degenerate: the three long diagonals)
{8/2}  = 2 x {4/1}    two squares
{9/3}  = 3 x {3/1}    three triangles
{10/4} = 2 x {5/2}    two pentagrams
{12/3} = 3 x {4/1}    three squares
{12/4} = 4 x {3/1}    four triangles
{12/6} = 6 x {2/1}    six digons
{16/4} = 4 x {4/1}    four squares
{18/6} = 6 x {3/1}    six triangles
{20/8} = 4 x {5/2}    four pentagrams
{24/9} = 3 x {8/3}    three octagrams
{24/10}= 2 x {12/5}   two dodecagrams
```

COROLLARY 58.5 (Attainable stroke counts). The set of stroke counts achievable on N points is exactly the set of proper divisors of N. Verified for N = 3..24: for every N, { gcd(N,k) : 1 <= k <= N−1 } equals the divisors of N other than N itself. **You cannot draw a five-stroke**

figure on twelve points with a constant step, because 5 does not divide 12, and no amount of ingenuity changes that.

COROLLARY 58.6 (Mirror symmetry). `gcd(N,k) = gcd(N, N−k)` for all `k`. Verified with zero counterexamples for `N = 3..100`. This halves the table: `k` beyond `N/2` retraces the figures of `k` below it, in reverse. It is why §58.5 tabulates `k = 1..floor(N/2)` and nothing is lost.

COROLLARY 58.7 (Vertex angle). The point angle of `{N/k}` is `180∗(N − 2k)/N` degrees, and the density (winding number) of each stroke is `k/g`, so each stroke's total turning is `360∗(k/g)` degrees. Verified: `{5/2}` gives `180∗1/5 = 36` degrees — the pentagram's point angle, which is the value every draughtsman knows.

58.3 Status table

Element	Status	Warrant
Strokes = `gcd(N,k)`, each of length `N/gcd(N,k)`	**PROVED** (group theory)	§58.1
Verified by exhaustive trace, `N = 3..40`, all k, 779 cases, 0 mismatches	**MEASURED**	§58.1
Unicursal iff `gcd(N,k) = 1`; count of unicursal `k` is `phi(N)`	**PROVED**	§58.2
Attainable stroke counts = proper divisors of `N`	**PROVED** (verified N=3..24)	§58.2
Pen lifts = `gcd(N,k) − 1`	**PROVED**	§58.2
Point angle `180(N−2k)/N` degrees	**PROVED** (geometry)	§58.2
The Spare reduction is a lossy, non-invertible encoding	**PROVED** (definition + count)	§58.12
Its output is bounded by 26 bits for any input length	**PROVED** (2^26 subsets)	§58.12
Measured letter-set collision rate on 218,002 words: 71.4%	**MEASURED**	§58.12
`gcd(N, .)` is a non-monotone, discontinuous function of k	**PROVED** (computed jumps)	§58.10
"The sigil is a nonlinear computing map"	**METAPHOR** (author's, badged)	§58.4
That drawing a figure produces any external effect	**REFUSED**	CANON §5

58.4 The metaphor, and exactly where it stops

The author's framing is that a sigil can be *read* as a nonlinear computing map. Here is the honest anatomy of that sentence.

What is literally true. The figure has a computational structure in the ordinary sense: a set of nodes, a deterministic successor rule, orbits, a component count, a path length per component,

and a total edge count. `gcd(N,k)` is precisely the number of *independent* closed processes the figure contains and `N/gcd(N,k)` is precisely the length of each. If you were scheduling work around a ring of N stations by handing it on k places at a time, those numbers would be the number of parallel workstreams and the length of each. **The number of separate strokes in the figure is `gcd(N,k)`, and that is the figure's parallelism** — that sentence is not a metaphor. It is a statement about a graph, and it is provable.

What is metaphor. Calling the figure a *computing map* imports the suggestion that something is being computed, by something, for some purpose. Nothing in the mathematics supplies an input, an output, or an agent. A star polygon is a static graph. It has the structural signature of a parallel process; it does not run one. **The metaphor is a way of seeing, it earns its keep by making the invariant visible, and it is not an assertion about the figure's nature.**

In what precise sense "nonlinear". The step map `x -> x + k (mod N)` is affine, hence linear in the technical sense — this must be conceded immediately, because it is the obvious objection. The nonlinearity is not in the map; it is in the **dependence of the figure's structure on the parameter k.** The function `k -> gcd(N,k)` is not monotone, not continuous, not smooth, and not approximable by any local rule: it is a number-theoretic function that jumps. On 24 points, k=11 gives one stroke and k=12 gives twelve. A unit change in the parameter changes the structure by a factor of twelve. §58.10 tabulates the maximum jump for every N in range. **That is the exact, defensible content of "nonlinear" here: the structure is a non-smooth arithmetic function of the design parameter, so you cannot tune it by small steps and you cannot interpolate.** Everything else in the phrase is metaphor and is badged as such.

What is refused, permanently. That drawing the figure does anything outside the person drawing it. Not asserted, not implied, not entertained.

The Engineering

58.5 The complete table, N = 3 to 24

Computed by explicit trace and cross-checked against `gcd`, with k running from 1 to `floor(N/2)`; by Corollary 58.6 the remaining k mirror these exactly. "Strokes" is the number of separate closed paths, "stroke length" the number of points (and edges) in each, and "compound form" the reduction of Corollary 58.4.

N	k	gcd(N,k)	strokes	stroke length N/g	unicursal	compound form
3	1	1	1	3	YES	–
4	1	1	1	4	YES	–
4	2	2	2	2	no	2 x {2/1}
5	1	1	1	5	YES	–
5	2	1	1	5	YES	–
6	1	1	1	6	YES	–
6	2	2	2	3	no	2 x {3/1}

6	3	3	3	2	no	3 x {2/1}
7	1	1	1	7	YES	–
7	2	1	1	7	YES	–
7	3	1	1	7	YES	–
8	1	1	1	8	YES	–
8	2	2	2	4	no	2 x {4/1}
8	3	1	1	8	YES	–
8	4	4	4	2	no	4 x {2/1}
9	1	1	1	9	YES	–
9	2	1	1	9	YES	–
9	3	3	3	3	no	3 x {3/1}
9	4	1	1	9	YES	–
10	1	1	1	10	YES	–
10	2	2	2	5	no	2 x {5/1}
10	3	1	1	10	YES	–
10	4	2	2	5	no	2 x {5/2}
10	5	5	5	2	no	5 x {2/1}
11	1	1	1	11	YES	–
11	2	1	1	11	YES	–
11	3	1	1	11	YES	–
11	4	1	1	11	YES	–
11	5	1	1	11	YES	–
12	1	1	1	12	YES	–
12	2	2	2	6	no	2 x {6/1}
12	3	3	3	4	no	3 x {4/1}
12	4	4	4	3	no	4 x {3/1}
12	5	1	1	12	YES	–
12	6	6	6	2	no	6 x {2/1}
13	1	1	1	13	YES	–
13	2	1	1	13	YES	–
13	3	1	1	13	YES	–
13	4	1	1	13	YES	–
13	5	1	1	13	YES	–
13	6	1	1	13	YES	–
14	1	1	1	14	YES	–
14	2	2	2	7	no	2 x {7/1}
14	3	1	1	14	YES	–
14	4	2	2	7	no	2 x {7/2}
14	5	1	1	14	YES	–
14	6	2	2	7	no	2 x {7/3}

14	7	7	7	2	no	7 x {2/1}
15	1	1	1	15	YES	–
15	2	1	1	15	YES	–
15	3	3	3	5	no	3 x {5/1}
15	4	1	1	15	YES	–
15	5	5	5	3	no	5 x {3/1}
15	6	3	3	5	no	3 x {5/2}
15	7	1	1	15	YES	–
16	1	1	1	16	YES	–
16	2	2	2	8	no	2 x {8/1}
16	3	1	1	16	YES	–
16	4	4	4	4	no	4 x {4/1}
16	5	1	1	16	YES	–
16	6	2	2	8	no	2 x {8/3}
16	7	1	1	16	YES	–
16	8	8	8	2	no	8 x {2/1}
17	1	1	1	17	YES	–
17	2	1	1	17	YES	–
17	3	1	1	17	YES	–
17	4	1	1	17	YES	–
17	5	1	1	17	YES	–
17	6	1	1	17	YES	–
17	7	1	1	17	YES	–
17	8	1	1	17	YES	–
18	1	1	1	18	YES	–
18	2	2	2	9	no	2 x {9/1}
18	3	3	3	6	no	3 x {6/1}
18	4	2	2	9	no	2 x {9/2}
18	5	1	1	18	YES	–
18	6	6	6	3	no	6 x {3/1}
18	7	1	1	18	YES	–
18	8	2	2	9	no	2 x {9/4}
18	9	9	9	2	no	9 x {2/1}
19	1	1	1	19	YES	–
19	2	1	1	19	YES	–
19	3	1	1	19	YES	–
19	4	1	1	19	YES	–
19	5	1	1	19	YES	–
19	6	1	1	19	YES	–
19	7	1	1	19	YES	–
19	8	1	1	19	YES	–
19	9	1	1	19	YES	–

20	1	1	1	20	YES	–
20	2	2	2	10	no	2 x {10/1}
20	3	1	1	20	YES	–
20	4	4	4	5	no	4 x {5/1}
20	5	5	5	4	no	5 x {4/1}
20	6	2	2	10	no	2 x {10/3}
20	7	1	1	20	YES	–
20	8	4	4	5	no	4 x {5/2}
20	9	1	1	20	YES	–
20	10	10	10	2	no	10 x {2/1}
21	1	1	1	21	YES	–
21	2	1	1	21	YES	–
21	3	3	3	7	no	3 x {7/1}
21	4	1	1	21	YES	–
21	5	1	1	21	YES	–
21	6	3	3	7	no	3 x {7/2}
21	7	7	7	3	no	7 x {3/1}
21	8	1	1	21	YES	–
21	9	3	3	7	no	3 x {7/3}
21	10	1	1	21	YES	–
22	1	1	1	22	YES	–
22	2	2	2	11	no	2 x {11/1}
22	3	1	1	22	YES	–
22	4	2	2	11	no	2 x {11/2}
22	5	1	1	22	YES	–
22	6	2	2	11	no	2 x {11/3}
22	7	1	1	22	YES	–
22	8	2	2	11	no	2 x {11/4}
22	9	1	1	22	YES	–
22	10	2	2	11	no	2 x {11/5}
22	11	11	11	2	no	11 x {2/1}
23	1	1	1	23	YES	–
23	2	1	1	23	YES	–
23	3	1	1	23	YES	–
23	4	1	1	23	YES	–
23	5	1	1	23	YES	–
23	6	1	1	23	YES	–
23	7	1	1	23	YES	–
23	8	1	1	23	YES	–
23	9	1	1	23	YES	–
23	10	1	1	23	YES	–
23	11	1	1	23	YES	–
24	1	1	1	24	YES	–
24	2	2	2	12	no	2 x {12/1}

```
 24    3         3         3                  8   no          3 x {8/1}
 24    4         4         4                  6   no          4 x {6/1}
 24    5         1         1                 24   YES          -
 24    6         6         6                  4   no          6 x {4/1}
 24    7         1         1                 24   YES          -
 24    8         8         8                  3   no          8 x {3/1}
 24    9         3         3                  8   no          3 x {8/3}
 24   10         2         2                 12   no         2 x {12/5}
 24   11         1         1                 24   YES          -
 24   12        12        12                  2   no         12 x {2/1}
---  ---  --------  -------  -----------------  --------  ------------
```

143 rows. 89 of them are unicursal — 62.2%. Every row was produced by walking the figure, not by evaluating `gcd`, and then checked against `gcd`.

58.6 The named verifications, with orbits printed

The theorem is worth distrusting until you have seen the walks. Here they are, as vertex sequences on Z_N starting from 0.

```
 {6/2}   gcd 2   TWO strokes of length 3:    [0,2,4]  and  [1,3,5]
 {6/3}   gcd 3   THREE strokes of length 2:  [0,3]  [1,4]  [2,5]
 {7/2}   gcd 1   ONE stroke of length 7:     [0,2,4,6,1,3,5]
 {7/3}   gcd 1   ONE stroke of length 7:     [0,3,6,2,5,1,4]
 {5/2}   gcd 1   ONE stroke of length 5:     [0,2,4,1,3]
 {8/3}   gcd 1   ONE stroke of length 8:     [0,3,6,1,4,7,2,5]
 {9/3}   gcd 3   THREE strokes of length 3:  [0,3,6]  [1,4,7]  [2,5,8]
 {10/3}  gcd 1   ONE stroke of length 10:    [0,3,6,9,2,5,8,1,4,7]
{12/3} gcd 3  THREE strokes of length 4: [0,3,6,9] [1,4,7,10] [2,5,8,11]
 {12/5}  gcd 1   ONE stroke of length 12:    [0,5,10,3,8,1,6,11,4,9,2,7]
{24/7} gcd 1  ONE stroke of length 24:   [0,7,14,21,4,11,18,1,8,15,22,5,12,19,2,9,16,2
                                          6,13,20,3,10,17]
```

The hexagram is literally two separate triangles: `[0,2,4]` and `[1,3,5]`, disjoint, sharing no vertex and no edge. The two triangles overlap visually on the page and are combinatorially independent. That is not a reading of the figure; it is what happens when you draw it. Every draughtsman who has ever drawn a Star of David knows they had to lift the pen once, and Corollary 58.3 says exactly why: `gcd(6,2) - 1 = 1` lift, always, no technique available.

The pentagram, by contrast, is `gcd(5,2) = 1` and needs zero lifts, which is why it is the figure that gets drawn in one motion in every schoolyard on earth.

58.7 The unicursal census

Per N, the number of unicursal steps and how the divisor structure of N controls it.

N	phi(N)	distinct proper star polygons	unicursal k in 1..floor(N/2)	total k	fraction	divisors of N
3	2	0	1	1	1.000	1, 3
4	2	0	1	2	0.500	1, 2, 4
5	4	1	2	2	1.000	1, 5
6	2	0	1	3	0.333	1, 2, 3, 6
7	6	2	3	3	1.000	1, 7
8	4	1	2	4	0.500	1, 2, 4, 8
9	6	2	3	4	0.750	1, 3, 9
10	4	1	2	5	0.400	1, 2, 5, 10
11	10	4	5	5	1.000	1, 11
12	4	1	2	6	0.333	1, 2, 3, 4, 6, 12
13	12	5	6	6	1.000	1, 13
14	6	2	3	7	0.429	1, 2, 7, 14
15	8	3	4	7	0.571	1, 3, 5, 15
16	8	3	4	8	0.500	1, 2, 4, 8, 16
17	16	7	8	8	1.000	1, 17
18	6	2	3	9	0.333	1, 2, 3, 6, 9, 18
19	18	8	9	9	1.000	1, 19
20	8	3	4	10	0.400	1, 2, 4, 5, 10, 20
21	12	5	6	10	0.600	1, 3, 7, 21
22	10	4	5	11	0.455	1, 2, 11, 22
23	22	10	11	11	1.000	1, 23
24	8	3	4	12	0.333	1, 2, 3, 4, 6, 8, 12, 24

Prime `N` gives a unicursal fraction of exactly 1.000: every step draws one continuous line. Highly composite `N` gives the worst ratios — `N = 12`, 18 and 24 all sit at 0.333. A designer who wants many available single-stroke figures should choose a prime number of points; a designer who wants many available *parallel decompositions* should choose a highly composite one. The two goals are exactly opposed, and the opposition is the divisor structure of `N`.

The column "distinct proper star polygons" counts `k` with `2 <= k < N/2` and `gcd(N,k) = 1` — the figures that are genuinely star-shaped, single-stroke, and not the convex polygon. For `N = 3, 4,` 6 it is zero: **there is no star polygon on three, four or six points, which is a fact about divisors and not about draughtsmanship.**

58.8 The parallelism reading

This section is where the invariant becomes an engineering quantity. Take the figure as a *schedule*: `N` stations around a ring, work handed on `k` places at a time, one hand-off per unit time.

```
P  = gcd(N,k)      number of independent workstreams (strokes)
W  = N             total work: one hand-off per station, always
T  = N / gcd(N,k)  critical path: the length of the longest stroke
speedup = W / T = gcd(N,k) = P
```

N	k	P = strokes	W = edges	T = critical path	speedup
6	2	2	6	3	2.0
12	3	3	12	4	3.0
12	4	4	12	3	4.0
24	8	8	24	3	8.0
20	5	5	20	4	5.0
7	2	1	7	7	1.0
13	5	1	13	13	1.0

Total work is `N` for every k; only the arrangement changes. That is the honest form of the parallelism claim: the figure never does less work, it does the same work in fewer sequential steps. The speedup is exactly `gcd(N,k)` and it is perfect — no overhead appears in the geometry, because the strokes share no vertex.

Now the constraint that makes this interesting rather than free. **The g strokes are *disjoint*. No point on one stroke is ever adjacent to a point on another.** In the vocabulary of Chapter 57 and Paper A, partitioning the ring into `g` strokes means that within one period, no pair from different strokes interacts at all. The cross-block term `u_cross` is unavailable inside the period. The surplus ceiling available to a stroke of `x = N/g` members is `ln x`, not `ln N` (Paper D, Corollary 2.2).

```
parallelism g:      latency falls from N to N/g
surplus ceiling:    falls from ln N to ln (N/g)
```

Parallelism buys latency and pays for it in ceiling. That trade is exact and it is the whole design decision.

58.9 The rate theorem: why stroke length three is optimal

Define the *ceiling rate* of a decomposition as surplus ceiling per unit of latency:

```
R(x) = ln(x) / x           where x = N/g is the stroke length
```

This is an upper bound on the rate at which a stroke can generate surplus, not an achieved rate; it is attained only if u can be held at 1 within each stroke. Badge it as a ceiling.

```
dR/dx = (1 - ln x) / x^2  =  0   at   ln x = 1,  i.e.  x = e = 2.718282
```

R rises on `(0, e)` and falls on `(e, infinity)`. The optimum is at `x = e`, and since stroke length must be an integer dividing `N`, the best available integer is 3 whenever 3 divides `N`. Computed:

```
x = 2   R = 0.346574
x = 3   R = 0.366204     <-- maximum over integers
x = 4   R = 0.346574
x = 5   R = 0.321888
```

```
x = 6   R = 0.298627
x = 12  R = 0.207076
```

Note the exact tie between `x = 2` and `x = 4`: `ln 2 / 2 = ln 4 / 4` because `ln 4 = 2 ln 2`. The tie is arithmetic, not coincidence.

For `N = 12`, sweeping the step:

k	g = gcd(12,k)	strokes	stroke length x	latency	ceiling ln x	rate ln(x)/x
1	1	1	12	12	2.484907	0.207076
2	2	2	6	6	1.791759	0.298627
3	3	3	4	4	1.386294	0.346574
4	4	4	3	3	1.098612	0.366204
5	1	1	12	12	2.484907	0.207076
6	6	6	2	2	0.693147	0.346574

The rate-optimal decomposition of twelve points is `{12/4}` — four triangles — and it is optimal because `ln x / x` peaks at `x = e` and 3 is the nearest integer divisor. The triad recurs across every organisational tradition this manual has touched, from Ostrom's nested units to the three-person review cell, and here it appears as the integer nearest to e under a log-over-linear objective. That is a derivation, offered as one, with no suggestion that anyone in history derived it.

The counterweight is Corollary 58.5 and the previous section: `{12/1}` and `{12/5}` are the only configurations that keep the full `ln 12 = 2.484907` ceiling, and they take twelve steps to get there. **Fast and shallow, or slow and deep. The figure makes you choose, and the divisors of `N` are the menu.**

58.10 In what sense the map is nonlinear

The concession first: the step map is affine. The nonlinearity claimed for this chapter lives in the parameter dependence, and here it is measured. For each `N`, the maximum change in stroke count caused by a single unit change in k:

N	max abs jump	at k -> k+1	strokes go
3	0	1 -> 2	1 -> 1
4	1	2 -> 3	2 -> 1
6	1	4 -> 5	2 -> 1
8	3	4 -> 5	4 -> 1
9	2	6 -> 7	3 -> 1
10	3	5 -> 6	5 -> 2
12	5	6 -> 7	6 -> 1
14	5	7 -> 8	7 -> 2
15	4	10 -> 11	5 -> 1
16	7	8 -> 9	8 -> 1
18	7	9 -> 10	9 -> 2

N	max abs jump	at k -> k+1	strokes go
20	9	10 -> 11	10 -> 1
21	6	13 -> 14	1 -> 7
22	9	11 -> 12	11 -> 2
24	11	12 -> 13	12 -> 1
5, 7, 11, 13, 17, 19, 23	0	—	always 1

On 24 points, moving the step from 12 to 13 takes the figure from twelve independent strokes to one, in a single unit of the parameter. No local model predicts that. A designer who tunes k by small increments and watches the result gets no gradient to follow, because there is no gradient: `gcd(N, .)` is constant-1 across every k coprime to N and spikes to a divisor exactly on the multiples.

Prime N is the degenerate case with maximum jump 0 — perfectly flat, always unicursal, no structure to tune at all. **The design lesson is exact: if you want the decomposition to be a tunable parameter, choose composite N; if you want it to be immune to a mis-set parameter, choose prime N.**

58.11 The unicursal hexagram, and how the tradition escaped the theorem

Here is a case where the historical record and the theorem meet head-on, and the record turns out to record the theorem's constraint.

{6/2} is two triangles and cannot be drawn without a lift. The nineteenth- and twentieth-century occult revival nonetheless produced a figure called the *unicursal hexagram*: a six-pointed star drawn in a single continuous line. This does not contradict Theorem 58.1, and the reason is worth stating precisely. **The theorem is about a CONSTANT step k. Relax that, and the obstruction disappears.**

Exhaustive search over all closed tours of the six points that use only non-adjacent chords — that is, only chords of length 2 or 3, the chords that make a figure look like a star rather than a hexagon — gives:

```
closed 6-cycles using only chords of length 2 or 3 :  6 vertex sequences
                                                      3 distinct edge sets
                                          all with the same chord multiset:
                                       four chords of length 2, two of length 3
```

The three edge sets are rotations of one another. **Up to symmetry there is exactly one such figure, its step sequence is `2, 2, 3, 2, 2, 3` rather than a constant, and it is the unicursal hexagram.** The tradition found, by drawing, the unique six-point star-chord Hamiltonian cycle. It did not have the theorem, so it could not have known that the figure was unique or why the regular hexagram resisted; it had the figure.

This is the clearest single illustration of what Part XI claims and of what it refuses. The claim: the old practice was operating on a real structural constraint and its artefacts record the constraint accurately. The refusal: nothing here says the figure does anything.

58.12 Spare's method is lossy compression, and here is its bit budget

Austin Osman Spare's sigil procedure, as described in his *The Book of Pleasure (Self-Love)* (1913), has three steps:

1. Write a statement of intent as a sentence.
2. Strike out every repeated letter, keeping first occurrences only.
3. Recombine the remaining letters into a single glyph.

Steps 1 and 2 constitute an encoding operation, and it may be described as one precisely. Step 2 is a map from strings to sets: it discards multiplicity and it discards order. The formal object is

```
sigma : strings over A  ->  subsets of A
sigma(w) = { letters occurring in w }
```

with the first-occurrence ordering retained only as a drawing convenience.

The hard ceiling. The codomain of `sigma` is the power set of a 26-letter alphabet.

```
|codomain| = 2^26 = 67,108,864
output entropy <= log2(2^26) = 26 bits, FOR ANY INPUT OF ANY LENGTH
```

That bound does not grow. A statement of intent of ten words or ten thousand produces at most 26 bits. Layer by layer:

retained letters m	C(26, m)	bits
8	1,562,275	20.575
10	5,311,735	22.341
12	9,657,700	23.203
13	10,400,600	23.310
14	9,657,700	23.203
16	5,311,735	22.341
18	1,562,275	20.575
20	230,230	17.813

The widest layer is `m = 13`, at 23.31 bits — a sigil that retains exactly half the alphabet carries the most information any sigil can carry, and it is under 24 bits.

Measured on real statements of intent, using a uniform-letter upper bound of `log2(26) = 4.700` bits per input letter:

statement	letters	retained	ratio	input bits (bound)	output bits	destroyed
THIS IS MY WILL TO BUILD A QUIET ROOM	29	16	0.552	136.31	22.34	113.97

statement	letters	retained	ratio	input bits (bound)	output bits	destroyed
IT IS MY WILL TO FINISH THE MANUAL	27	14	0.519	126.91	23.20	103.71
MY WILL IS TO SLEEP WITHOUT FEAR	26	15	0.577	122.21	22.88	99.33
THIS IS MY WISH TO OWN A QUIET ROOM	27	14	0.519	126.91	23.20	103.71
I WILL BUILD THE INSTRU-MENT AND KEEP IT	32	16	0.500	150.41	22.34	128.07

Retained sets, in first-occurrence order:

```
THIS IS MY WILL TO BUILD A QUIET ROOM     ->  THISMYWLOBUDAQER
IT IS MY WILL TO FINISH THE MANUAL        ->  ITSMYWLOFNHEAU
MY WILL IS TO SLEEP WITHOUT FEAR          ->  MYWILSTOEPHUFAR
THIS IS MY WISH TO OWN A QUIET ROOM       ->  THISMYWONAQUER
I WILL BUILD THE INSTRUMENT AND KEEP IT   ->  IWLBUDTHENSRMAKP
```

Preimage size — the measure of how thoroughly the operation is not invertible. The number of length-n strings that use exactly a given m-letter set is a surjection count, `sum over i of (-1)^i C(m,i) (m-i)^n`. Computed:

```
length 28 over a given 15-letter set:   4.716e+31   = 2^105.2 strings
length 30 over a given 14-letter set:   3.691e+33   = 2^111.5
length 32 over a given 16-letter set:   2.499e+37   = 2^124.2
length 36 over a given 17-letter set:   1.787e+43   = 2^143.7
```

Roughly 10^31 to 10^43 sentences of the right length collapse onto any one sigil's letter set. The operation is a digest, not a code: it cannot be decoded, and that is by construction, not by accident.

Measured on a natural-language corpus for calibration: over the 218,002 alphabetic words of length 6 or more in `/usr/share/dict/words`, there are 101,517 distinct letter-sets, mean class size 2.147, and **71.4% of words share their letter-set with at least one other word.** Input entropy 17.734 bits, output entropy 16.037 bits, 1.697 bits destroyed — far gentler than the sentence

case, because words are short and the reduction has less multiplicity to throw away. The larger the input, the more brutal the compression, and a sentence is a large input.

58.13 The attention-binding claim, badged and bounded

The tradition attaches a second step to the encoding: the operator is to charge the glyph and then forget the sentence, retaining the form and losing the wording. **This manual can describe the encoding exactly and can describe the second step only as a claim about the operator's own perceptual and attentional state — and it makes no claim of any effect outside the operator.**

What can be said in the formalism, and stated with its badge on:

(a) The reduction genuinely makes the source unrecoverable, and that is measured. The 10^31-to-1 preimage figure above is not a metaphor for forgetting; it is the actual non-invertibility of the map. If the practitioner's stated aim is to hold a form while the wording becomes irrecoverable, the operation performs exactly that, as arithmetic.

(b) A glyph is a small, stable, repeatable perceptual object. In the vocabulary of Chapter 34, `M = (D/P) E`, a glyph is an environment term `E` with very low variance across presentations. Chapter 59 develops the consequence: a low-variance, repeatedly presented `E` is *periodic forcing* of a leaky integrator, and cadence beats intensity for that reason and no other.

(c) The compression is the point, mnemonically. Chapter 57 measured the cost of holding relations: `n(n−1)/2` unstructured versus `n` structured. A sigil is the same move at the extreme — one object standing for a proposition — and the reason it is cheap to rehearse is that it is under 24 bits.

(d) Nothing above is evidence that the practice achieves its traditional aims. No efficacy study is cited here, because this manual will not fabricate one. What is offered is a description of the encoding and an honest statement of what the formalism can and cannot reach. The formalism reaches the operator's own storage law — `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`, Chapter 40 — and it stops at the operator's skin. **Whether an intention held with high `eta` and low `delta` in one person's stored order changes that person's subsequent behaviour is an empirical question with an obvious affirmative answer and no mystery in it; whether a drawn figure does anything else is a question this manual refuses.**

58.14 Composing the two halves of the chapter

The two results join. Lay the retained letters of a Spare reduction on the vertices of a regular `N`-gon, `N` = the number of retained letters, and connect them with a constant step `k`. Then:

```
N        = number of retained letters (measured: 14 to 17 in the examples above)
  gcd(N, k)    = number of separate strokes the glyph requires
  N / gcd(N,k) = letters bound into each stroke
  unicursal    <=>  gcd(N, k) = 1
```

For `N = 14` (from `IT IS MY WILL TO FINISH THE MANUAL`), the table in §58.5 gives: `k = 1, 3, 5` unicursal; `k = 2, 4, 6` give two strokes of seven; `k = 7` gives seven strokes of two. `phi(14) = 6`, so 6 of the 13 non-trivial steps draw a single line.

For `N = 16` (from `THIS IS MY WILL TO BUILD A QUIET ROOM`), `k = 1, 3, 5, 7` unicursal; `k = 4` gives four squares; `k = 8` gives eight digons. `phi(16) = 8`.

The composition is a design rule, not a discovery: if you want one stroke, choose a step coprime to the letter count; if you want the glyph to fall into `g` visually separable parts, choose `k` a multiple of a divisor `g` of `N`, and the parts will each carry `N/g` letters. It is exactly the same rule that governs the review-cycle schedule in the worked example below, which is the point of the chapter and the point of Part XI.

The Specification

S58.1 — Compute `gcd(N,k)` before drawing, scheduling, or committing to a figure. It is one line of arithmetic and it determines the stroke count, the stroke length, the pen lifts and the parallelism. Print it.

S58.2 — State `N` and `k` explicitly in any figure specification. A figure described as "a six-pointed star" is under-specified: `{6/2}` and the unicursal hexagram are different objects with different stroke counts.

S58.3 — Choose prime `N` for guaranteed unicursality; choose highly composite `N` for a rich decomposition menu. `phi(N)/N` is the unicursal fraction and the divisor set is the menu; both are printed in §58.7 for `N = 3..24`. These goals are opposed and the choice must be made once, early, and on purpose.

S58.4 — Never expect to tune `k` by small steps. `gcd(N, .)` has no gradient. Enumerate the divisors, evaluate each candidate, and pick. §58.10 shows single-unit jumps as large as 11.

S58.5 — Report total work separately from latency. `W = N` always; `T = N/g`; `speedup = g`. A parallelisation that appears to reduce work has been mis-measured.

S58.6 — Price the ceiling loss whenever you buy parallelism. Splitting into `g` disjoint strokes lowers the per-period surplus ceiling from `ln N` to `ln(N/g)`. Print both numbers side by side before deciding.

S58.7 — Use the rate criterion where throughput matters. `R(x) = ln x / x`, maximised at `x = e`; the best integer stroke length is 3 when 3 divides `N`. Badge it as a ceiling rate, not an achieved rate.

S58.8 — Use the full-ceiling configuration where depth matters. `gcd(N,k) = 1` is the only way to keep `ln N`, and it costs `N` steps of latency. Depth and speed are not simultaneously available.

S58.9 — When a figure must be unicursal and `gcd` forbids it, vary the step. The obstruction is a theorem about constant steps only. Enumerate the Hamiltonian cycles under whatever chord constraint you actually need — §58.11 does this for six points and finds exactly one figure up to symmetry.

S58.10 — Describe the Spare reduction as an encoding, in those words. It is a map from strings to subsets. Report the retained-letter count `m`, the output bound `log2 C(26,m)` bits, and the preimage magnitude. Do not describe it as compression without saying *lossy*.

S58.11 — State the 26-bit ceiling whenever sigil information content is discussed. Output entropy is bounded by 26 bits for any input length; the widest layer is 23.31 bits at `m = 13`.

S58.12 — Keep the metaphor badge on every occurrence. "Computing map" is the author's metaphor and must be marked at each use. The parallelism statement — strokes = `gcd(N,k)` — is not a metaphor and should be stated without hedging, because it is a theorem.

S58.13 — Claim nothing outside the operator. Any statement about a sigil's effect is bounded by the operator's own storage law and stops at the operator's skin. See §58.13(d).

S58.14 — Verify by walking, not by formula, the first time. Every table in this chapter was produced by explicit trace and then checked against `gcd`. A theorem you have not tested on the object is a theorem you have not understood.

The Failure Modes

F58.1 — SILENT LIFT. *Signature:* a figure or a schedule is specified as continuous, and in execution someone lifts the pen or breaks the chain, and no one records it. *Mechanism:* `gcd(N,k) > 1`. The break was structurally required, not a lapse. *Fix:* compute `gcd` first. If `g > 1`, declare the `g − 1` breaks as part of the specification and put a named owner on each. An undeclared break becomes an untracked hand-off, which is the classic Conway's Law seam.

F58.2 — PHANTOM PARALLELISM. *Signature:* a schedule is described as having `g` parallel workstreams, and work in fact crosses between them constantly. *Mechanism:* the strokes of `{N/k}` are *disjoint*. If real traffic crosses, the model is not `{N/k}`, and `gcd` does not describe the actual topology. *Fix:* measure the cross-stroke traffic. If it is non-zero, the correct model is a graph with those edges added, and the component count must be recomputed on the real graph, not the idealised one.

F58.3 — CEILING BLINDNESS. *Signature:* a team is split into small parallel cells for speed, throughput improves, and the quality of what emerges degrades in a way no one can locate. *Mechanism:* the ceiling fell from `ln N` to `ln(N/g)`. For `N = 12, g = 4`: from 2.484907 to 1.098612 nats, a 55.8% reduction in the maximum surplus any single period can produce. *Fix:* S58.6. Print both ceilings. If depth is the requirement, run at least some periods at `gcd = 1`; alternate configurations rather than committing to one.

F58.4 — GRADIENT HUNTING. *Signature:* a parameter is being tuned by small increments and the behaviour is erratic, so the increments get smaller. *Mechanism:* the response is `gcd(N, k)`, which is discontinuous everywhere it matters. Smaller steps do not help; they cannot. *Fix:* stop tuning. Enumerate the divisors of `N`, evaluate the finite candidate set exhaustively, and choose. There are at most `floor(N/2)` candidates and for `N <= 24` that is at most 12.

F58.5 — GCD MYSTICISM. *Signature:* the arithmetic invariant is presented as though it explained the historical or symbolic *significance* of a figure, rather than its stroke count. *Mechanism:* an ordinary category error — a structural fact recruited to carry interpretive weight it cannot bear. It is the same error as F57.2 (hash mysticism), in a chapter with a better theorem and therefore a stronger temptation. *Fix:* the theorem says how many strokes. It says nothing else. Where a sentence about the figure would be false if the figure had been invented yesterday to organise a warehouse, that sentence is outside the theorem.

F58.6 — CONSTANT-STEP ASSUMPTION. *Signature:* "this figure cannot be drawn in one line" is asserted, and someone then draws it in one line. *Mechanism:* Theorem 58.1 governs constant steps. Variable-step tours are a different and larger search space. The unicursal hexagram lives in

the gap. *Fix:* state the step regime with the claim. “No constant step draws {6/2} in one line” is true; “no single line touches six points in a star pattern” is false, with exactly one counterexample up to symmetry.

F58.7 — DIGEST-AS-CODE. *Signature:* someone attempts to reconstruct the original statement of intent from a sigil, or treats the glyph as carrying the sentence. *Mechanism:* the preimage of a 15-letter set among 28-letter strings is about 2^105. The map is not invertible and the loss is total. *Fix:* keep the statement separately if you need it. If the practice requires that the statement be irrecoverable, the encoding already guarantees it; nothing further is needed and nothing can undo it.

F58.8 — UNBOUNDED-SIGIL FALLACY. *Signature:* a longer, more elaborate statement of intent is believed to produce a richer glyph. *Mechanism:* output entropy is capped at 26 bits regardless of input length, and at 23.31 bits in practice. Beyond about 13 retained letters the layer count *falls*: C(26,20) < C(26,13) by a factor of 45. *Fix:* if information content is the goal, target m near 13. If it is not the goal, do not invoke information content as a reason for elaboration.

F58.9 — PRIME COMPLACENCY. *Signature:* a prime N is chosen for its guaranteed unicursality, and later someone needs a parallel decomposition and cannot have one. *Mechanism:* prime N has no proper divisors, so gcd(N,k) = 1 for every k and the only attainable stroke count is 1. Maximum jump 0: perfectly robust, perfectly inflexible. *Fix:* decide early which property is load-bearing. If both may be needed, choose N with a small number of divisors — N = 2p for prime p gives exactly the stroke counts 1, 2, and p, which is a usable three-item menu with phi(2p) = p−1 unicursal steps.

Worked Example

A twelve-station design review, scheduled as a star polygon

A hardware group runs a design review across **twelve stations** — twelve specialists who must each see and sign every design package. The current process is a single chain: station 1 to 2 to 3, and so on to 12. One package per hand-off, one hand-off per working day. The group wants throughput and does not want to lose depth. We use the chapter.

Step 1 — Identify N and enumerate the menu (S58.1, S58.3).

```
N = 12          divisors: 1, 2, 3, 4, 6, 12
attainable stroke counts (proper divisors): 1, 2, 3, 4, 6
candidate steps k = 1..6 (Corollary 58.6 mirrors the rest)
phi(12) = 4, so only k = 1, 5, 7, 11 are unicursal: a 33.3% unicursal fraction
```

Twelve is the worst-case cardinality for single-stroke figures and the best case for decompositions. That is the trade the group has already made by having twelve people.

Step 2 — The configuration table. From §58.5 and §58.9, with the ceiling and rate columns:

Config	g = strokes	stroke size x	latency (days)	ceiling `ln x`	rate `ln x / x`
{12/1} current	1	12	12	2.484907	0.207076
{12/2}	2	6	6	1.791759	0.298627
{12/3}	3	4	4	1.386294	0.346574
{12/4}	4	3	3	1.098612	0.366204
{12/5}	1	12	12	2.484907	0.207076
{12/6}	6	2	2	0.693147	0.346574

Explicit orbits for the leading candidate, {12/3}, stations numbered 0-11:

```
stroke A:  0 -> 3 -> 6 -> 9  -> 0
stroke B:  1 -> 4 -> 7 -> 10 -> 1
stroke C:  2 -> 5 -> 8 -> 11 -> 2
```

Three closed loops of four, disjoint, twelve hand-offs total, two pen lifts — that is, two declared seams between loops that require an explicit owner (F58.1).

Step 3 — Storage law with real parameters. The group's measured parameters:

```
eta   = 0.45    fraction of review interaction that becomes durable design knowledge
delta = 0.08    per-period decay (rotation onto other projects, standards drift)
 Q_t   = 6.0       current stored design capability, nat-scaled
C_t  = 0.05*(g-1) + 0.15   coordination cost: a base of 0.15 plus 0.05 per merge seam
```

The cost model is the group's own estimate, badged as an estimate: each additional stroke adds one merge point where the loops' conclusions must be reconciled, and the group prices that at 0.05.

One-step change, `dQ = eta * S - delta * Q_t - C_t`, taking S at each configuration's ceiling (the optimistic case; real S is lower and the ranking is unaffected):

Config	strokes	members/stroke	S ceiling	C_t	eta*S	dQ
{12/1}	1	12	2.484907	0.1500	1.1182	**+0.488208**
{12/2}	2	6	1.791759	0.2000	0.8063	+0.126292
{12/3}	3	4	1.386294	0.2500	0.6238	-0.106168
{12/4}	4	3	1.098612	0.3000	0.4944	-0.285624
{12/6}	6	2	0.693147	0.4000	0.3119	-0.568084

The rate-optimal figure {12/4} — four triangles, the configuration §58.9 identified as maximal in `ln x / x` — has a *negative* one-step change in stored order. The two results are not in conflict and the reason is exactly the trade in §58.8: {12/4} maximises surplus per unit of *latency*, and the storage law is denominated per *period*, not per unit of latency. If the group runs four three-person loops, each period produces at most 1.098612 nats of surplus, `eta` converts 0.494 of it, and decay plus merge cost consumes 0.780. Order falls.

The reconciliation is quantitative and it is the real finding. {12/4} completes a cycle in 3 days rather than 12 — four times as many periods per quarter. Over a 60-working-day quarter:

```
  {12/1}: 60/12 = 5 cycles   per-cycle dQ = +0.488208  quarter total = +2.441040
   {12/3}:  60/4  = 15 cycles   per-cycle dQ = -0.106168   quarter total = -
1.592513
   {12/4}:  60/3  = 20 cycles   per-cycle dQ = -0.285624   quarter total = -
5.712489
```

The single chain wins outright on this cost model, because the merge cost grows with stroke count while the ceiling falls with it, and running the loss more often only compounds the loss. The rate criterion `ln x / x` assumed the only cost was latency. It is not; `C_t` rises with `g` and `delta * Q_t` is charged every period regardless. **The rate theorem is a ceiling result and this is what it looks like when the ceiling is not the binding constraint.**

Step 4 — Headroom for each configuration. The master gauge, at the unpartitioned ceiling `ln 12`:

```
h = eta * ln 12 - delta * Q_t - C_t = 1.118208 - 0.480000 - C_t

{12/1}: h = +0.488208
{12/2}: h = +0.438208
{12/3}: h = +0.388208
{12/4}: h = +0.338208
{12/6}: h = +0.238208
```

All positive: the group has room to grow in every configuration *if* it can reach `ln 12`. Only `{12/1}` and `{12/5}` can. **Headroom exists and only the unicursal configurations can spend it, which is the entire decision in one line.**

Step 5 — The chosen design. The group adopts `{12/5}` rather than `{12/1}`, and the reason is not in the surplus arithmetic at all — the two are identical there, both `g = 1`, both twelve steps, both ceiling `ln 12`. The reason is the *sequence*:

```
{12/1}:  0,1,2,3,4,5,6,7,8,9,10,11
{12/5}:  0,5,10,3,8,1,6,11,4,9,2,7
```

`{12/1}` visits adjacent stations, which in this group means adjacent specialisms; `{12/5}` visits stations five apart, so each hand-off crosses to a distant specialism and returns. **Both have the same `N`, the same `g`, the same ceiling and the same latency; they differ in how much of the ceiling is reachable, because u between adjacent specialists is lower than u between distant ones.** That is Chapter 9's engineering of u reappearing as a choice of step, and the star polygon is the object that makes the choice enumerable: there are exactly `phi(12)/2 = 2` distinct unicursal figures on twelve points, `{12/1}` and `{12/5}`, and the group can simply compare them.

Estimating u from the group's own records — mean non-parallelism 0.31 for adjacent pairs, 0.58 for pairs five apart:

```
{12/1}: N_eff = 1 + 11*0.31 = 4.410000   F = ln 4.410000 = 1.483875 nats  (59.7% of ln 12)
{12/5}: N_eff = 1 + 11*0.58 = 7.380000   F = ln 7.380000 = 1.998774 nats  (80.4% of ln 12)

 dQ {12/1} = 0.45*1.483875 - 0.48 - 0.15 = +0.037744
 dQ {12/5} = 0.45*1.998774 - 0.48 - 0.15 = +0.269448
```

Changing the step from 1 to 5 — the same twelve people, the same twelve hand-offs, the same twelve days, one number changed in the schedule — multiplies the per-cycle gain in

stored order by 7.14. Over five cycles a quarter: +0.188718 against +1.347241.

That is the chapter's engineering claim in one worked number, and it rests on nothing beyond Theorem 58.1, Paper A, and Paper E. The figure told the group how many strokes were available and which steps were unicursal. The surplus law told them which unicursal step was worth taking. The storage law told them what it was worth.

A star polygon is a schedule with an exact arithmetic invariant, and reading it as one is not a metaphor. Reading it as a computing map is a metaphor, and the badge stays on. Chapter 59 takes the third artefact of the old notation — the memory palace and the ritual calendar — and finds the storage law waiting in both.

Chapter 59 . Memory Palace, Ritual, and Periodic Forcing

Chapter 57 read the correspondence table as a finite relational graph. Chapter 58 read the sigil as an encoding with an exact arithmetic invariant. This chapter takes the two remaining technologies of the old notation that are large enough to have architecture — **the memory palace and the liturgical calendar** — and shows that each is a named instance of a law already proved in this manual.

Neither reading requires anything to be added to the formalism. The memory palace is the storage law with space as the retention medium. The ritual calendar is a periodic forcing function applied to a first-order low-pass filter. Both statements are exact, both are computable in advance, and neither carries any claim beyond arithmetic.

Three results, named up front:

- **The loci method is not a trick for holding more. It is a technique for holding the same amount with a smaller `delta` and a larger `eta`, and the gain is the product of three separate factors, each measurable.** In the worked case it is a factor of 14.
- **The case for regular practice over heroic effort has two halves, and only one of them is the ceiling argument of Chapter 2.** The second half is Theorem 59.1 below, which follows from Paper A's strict concavity and has nothing to do with saturation: dividing a fixed budget of interaction into more, smaller occasions never delivers less, and usually delivers much more.
- **The comparison between one large impulse and the same total delivered periodically is not ambiguous, but it is subtler than the folklore.** Over infinite time and with no ceiling the two are exactly equal in integrated effect. Three things differ, and all three are computed here: the steady state, the minimum, and — once Paper D's Corollary 2.2 is applied — the deliverable total.

The chapter's load-bearing claim is the last one, and it is arithmetic rather than exhortation: **a single impulse, however large, has a steady state of exactly zero. A cadence has a steady state of `m/delta`. That is the whole difference, and everything else in the practice literature is a gloss on it.**

The Law

59.1 Statement one — the loci law

STATEMENT (the method of loci, read formally). Let a set of `N` items be encoded by binding each to a distinct location in a familiar, mentally traversable space. Then, in the storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`:

- the **non-parallelism `u`** rises, because two items bound to two different places are, by construction, framed differently;
- the **conversion efficiency `eta`** rises, because retrieval has an addressable index — the walk — rather than requiring undirected search;
- the **retention loss `delta`** falls, because the index is a spatial layout that the retriever already maintains for other reasons and does not have to keep alive on this material's account.

This is a reading, not a theorem. What is a theorem is what follows once the three parameter movements are granted: the steady state moves by their product,

```
Q_inf = eta F(u) / delta          (Paper E, Thm 3.5, with c* = 0)

Q_inf'   eta'   F(u')   delta
------ = ---- . ----- . ------
Q_inf    eta    F(u)    delta'
```

and the three factors are separately measurable, separately attackable, and multiply. §59.6 computes them.

The tradition is old, continuous, and explicit about its own mechanism. Cicero (*De Oratore* II.86) attributes the technique to Simonides of Ceos, said to have identified the dead at a collapsed banquet hall by recalling where each guest had been sitting. Whatever the anecdote's historicity, its content is the claim being made here: **the retrieval index was the seating plan.** The fullest surviving classical treatment is the anonymous *Rhetorica ad Herennium* (c. 86-82 BCE), Book III; Quintilian gives a second account in the *Institutio Oratoria*, Book XI. Frances Yates's *The Art of Memory* (1966) is the standard history of how the technique travelled from Roman rhetoric through the medieval and Renaissance traditions.

59.2 Statement two — the cadence theorem

Chapter 2 established that cadence beats intensity because surplus saturates at `ln N` per period while cost does not saturate. That argument depends entirely on the ceiling. **There is a second, independent argument that does not, and it is a theorem.**

THEOREM 59.1 (the cadence theorem). Let a fixed budget `T` of interaction be spread over `n` periods, `T/n` in each, and let the surplus generated in a period be `g(L)` where `L = T/n` is that period's allocation. If `g` is concave on `[0, infinity)` with `g(0) = 0`, then

```
D(n) := n * g(T/n)
```

is nondecreasing in `n`, and strictly increasing wherever `g` is strictly concave.

Proof. Treat `n` as a positive real. Write `L = T/n`, so `D = T g(L)/L`. Then

```
dD/dn = g(L) − L g'(L)
```

by the chain rule. Concavity of `g` with `g(0) = 0` gives `g(L) = g(L) − g(0) = integral_0^L g'(s) ds >= L g'(L)`, because `g'` is nonincreasing. Hence `dD/dn >= 0`, with strict inequality wherever `g'` is strictly decreasing. []

The hypothesis is satisfied by this manual's own surplus functional. Paper A, Theorem 3.3 proves that `f(u) = ln(1 + (N−1)u)` is strictly concave on `[0,1]`, and axiom S2 gives `f(0) = 0`. If the non-parallelism achieved in a session is any nondecreasing concave function of session length — and in particular if it is *linear* in length up to saturation, the simplest possible assumption — then the composite `g(L) = f(u(L))` is concave through the origin and Theorem 59.1 applies without further assumption.

Therefore the case for regular practice does not rest on the ceiling. It rests on the same concavity that makes the surplus law a logarithm, which is the same concavity Fechner found

in sensation. Diminishing returns within one occasion is precisely what makes many occasions worth more than one.

59.3 What is proven and what is postulated

Element	Status	Warrant
`D(n) = n g(T/n)` nondecreasing in `n` for concave `g`, `g(0)=0`	**PROVED**	§59.2
`f(u) = ln(1+(N−1)u)` is strictly concave with `f(0)=0`	**PROVED**	Paper A, Thm 3.3 and S2
A single impulse has steady state `Q_inf = 0`; a rate `m` has `Q_inf = m/delta`	**PROVED**	Paper E, Thm 3.5
Impulse and cadence have equal *integrated* effect absent a ceiling	**PROVED** (LTI)	§59.8, computed
Steady-state cycle: `Q_max = A/(1−a^P)`, `Q_min = Q_max a^P`	**PROVED**	Chapter 2, §2.13
Required retention for a rite of interval `P`: `a^P >= r/(1+r)`, `r = theta/A`	**PROVED** (algebra)	§59.14
Costly-requirement break-even: `dc <= (eta F − c)(1 − delta'/delta)`	**PROVED** (algebra)	§59.15
That the method of loci raises `u` and `eta` and lowers `delta`	**POSTULATE** (reading)	§59.1, badged
That a ritual calendar is well described by the storage law at all	**POSTULATE** (modelling)	CANON §1.4, badged
That any particular `(eta, delta, u)` triple describes a particular practice	**REQUIRES MEASUREMENT**	Chapters 3, 4, 9
Durkheim's, Rappaport's and Whitehouse's mechanisms for what ritual *does*	**SOCIAL THEORY**, not proved here	§59.13-59.14

59.4 Honesty clauses

- **Nothing in this chapter claims that a symbol, a place, a gesture or a date has any effect other than through a human being who encodes, retrieves, or attends.** The memory palace is a filing system. The calendar is a schedule. Both are effective; neither is mysterious.
- The illustrative parameter values in §59.6 and §59.8 are *design values*, chosen to make the structure legible, and they are labelled as such every time they appear. They are not

measurements of any population. The protocol for obtaining real ones is Chapter 3 (`delta`), Chapter 4 (`eta`), and Chapter 9 (u).

- As always, surplus as deployed is a **lower bound** on the ensemble's information (Papers A and D), so every `Q_inf` here is a floor.
- The anthropological literature cited in §59.13-59.17 is theory and observation about human societies. It is used here to name phenomena the formalism can represent. **The formalism does not validate the anthropology, and the anthropology does not validate the formalism.**
- Theorem 59.1 assumes g concave. If per-occasion response is genuinely *convex* over some range — if there is a warm-up cost below which nothing happens at all — the theorem reverses on that range, and §59.12 shows exactly how a fixed per-occasion cost restores an interior optimum. **The theorem does not say "shorter is always better." It says concavity favours division, and it names the term that stops the division.**

The Engineering

59.5 The *ad Herennium* rules are a parameter list

The classical treatment of artificial memory is unusually specific about procedure. The *Rhetorica ad Herennium* (III.16-24) gives rules for choosing loci and for constructing images. Read against the storage law, **each rule moves exactly one parameter, and the set of rules covers the parameter set without a remainder.**

Classical rule (*ad Herennium* III, paraphrased)	Parameter moved	Mechanism in the formalism
Choose loci in a deserted place, not a crowded one	`C` down	fewer competing encodings; less interference cost per period
Set the loci at moderate intervals, distinguishable one from the next	`u_cross` up	between-place non-parallelism is the block term `beta^2` of Thm 4.1
Do not let the loci be too like one another	`u_cross` up, `N` preserved	near-identical loci collapse to one effective locus; `N_eff` falls
Neither too bright nor too dark	`N` preserved	discriminability sets how many categories are actually distinct
Learn the series in order, and be able to begin from any point	`eta` up	multiple entry points make retrieval a graph traversal, not a list scan
Make the images striking, active, unusual	`u_div` up	within-place non-parallelism is the `alpha^2 sin^2 gamma` term
Use images of things you already care about	`eta` up	conversion into durable structure is higher for material with existing attachment

Classical rule (*ad Herennium* III, paraphrased)	Parameter moved	Mechanism in the formalism
Revisit the palace to keep the loci fresh	`delta` down	maintenance of the index rather than of the contents

That table is the chapter's method in miniature. No rule requires a mechanism outside the law, and no parameter of the law is left without a rule. The tradition arrived at a complete parameter-tuning procedure by iteration over roughly two thousand years and wrote it down in the vocabulary it had.

59.6 The arithmetic of place

Take a corpus of `N = 40` items. Paper D's Corollary 2.2 caps surplus at `ln 40 = 3.6889` nats per occasion, and no encoding exceeds it. What differs between encodings is how much of that ceiling they reach, and the driver is `u`.

Design values, illustrative and badged as such:

Encoding	`u`	`N_eff = 1+(N−1)u`	`F = ln N_eff`	% of `ln 40`
rote list, single undifferentiated context	0.15	6.850	1.9242	52.2%
rote plus grouping into chunks	0.40	16.600	2.8094	76.2%
loci, ordinary domestic places	0.75	30.250	3.4095	92.4%
loci, deliberately distinct places, vivid images	0.90	36.100	3.5863	97.2%
loci, idealised	0.97	38.830	3.6592	99.2%

The interesting feature of that column is how quickly it saturates: **an ordinary, undesigned palace already reaches 92% of the ceiling, and all the effort of the classical rules buys the last 8%.** That is the logarithm again. It is also the honest answer to why beginners get most of the benefit immediately and experts spend years on the remainder.

Palace size obeys an exact law and it is unforgiving. The ceiling is `ln L` in the number of loci `L`, so:

loci `L`	5	10	20	40	80	160	320
`ln L`	1.6094	2.3026	2.9957	3.6889	4.3820	5.0752	5.7683

Every doubling of the palace buys exactly `ln 2 = 0.693147` nats, forever. Not more at the start and not less at the end. Going from 5 loci to 10 buys the same as going from 160 to 320.

This is the most useful single fact for anyone designing a palace: **capacity is bought in fixed increments at exponentially rising cost in construction effort, so the question is never "how big" but "how many doublings can I afford to build and maintain."**

The two-source decomposition prices place and image separately. Paper A's Theorem 4.1 splits `u = u_cross + u_div` — between-locus and within-locus non-parallelism — and bounds the result:

```
max(f(u_cross), f(u_div))  <=  f(u)  <=  f(u_cross) + f(u_div)
```

Computed at `N = 40`:

u_cross (places)	u_div (images)	u	f(u_cross)	f(u_div)	lower	f(u)	upper	slack
0.85	0.05	0.90	3.5308	1.0818	3.5308	3.5863	4.6126	1.0263
0.80	0.10	0.90	3.4720	1.5892	3.4720	3.5863	5.0612	1.4749
0.45	0.45	0.90	2.9205	2.9205	2.9205	3.5863	5.8409	2.2546
0.60	0.20	0.80	3.1946	2.1748	3.1946	3.4720	5.3693	1.8974

Two readings, both practical. First, **`f(u)` sits very close to the lower bound in every row**, which is Theorem 3.4's subadditivity doing its work: the two sources overlap heavily and you do not get to add them. Second, **the balanced row (0.45/0.45) has the largest gap between the naive sum and the truth** — 2.25 nats of imagined surplus that does not exist. An operator who has invested equally in place-distinctness and image-vividness will overestimate the combined effect by more than one who has invested lopsidedly. That is a counterintuitive and checkable consequence.

The three-factor gain, computed. Period one day. Design values, badged.

Regime	u	eta	delta/day	F	eta F	Q_inf	DC gain eta/delta	tau	t_half
rote rehearsal	0.15	0.20	0.25	1.9242	0.3848	1.5394	0.800	3.48 d	2.41 d
rote + grouping	0.40	0.24	0.20	2.8094	0.6743	3.3713	1.200	4.48 d	3.11 d
loci encoding	0.90	0.38	0.06	3.5863	1.3628	22.7132	6.333	16.16 d	11.20 d
loci + spaced review	0.90	0.38	0.03	3.5863	1.3628	45.4264	12.667	32.83 d	22.76 d

Decomposing the loci-versus-rote ratio:

```
surplus factor    F2/F1        = 1.8637
conversion factor eta2/eta1    = 1.9000
retention factor  delta1/delta2 = 4.1667
--------------------------------------
```

```
product                        = 14.7546
direct Q_inf ratio             = 14.7546     (identical, as it must be)
```

Two thirds of the advantage is in `delta`. The surplus and conversion factors together give `1.8637 x 1.9000 = 3.54`; retention alone gives 4.17. **The memory palace is primarily a retention device and only secondarily an encoding device, and every popular account has that emphasis backwards.** Chapter 3's general rule reappears in miniature: retention is a gain term, and nobody treats it as one.

59.7 What the evidence supports, and what it does not

The method of loci is one of the few techniques in this manual with a genuine experimental literature, and the honest summary is narrower than the enthusiasm around it.

What is well supported. That the technique substantially improves ordered recall of arbitrary lists is not seriously contested; Bower's review in *American Scientist* (1970) is an early and clear statement of the experimental case. Maguire, Valentine, Wilding and Kapur (*Nature Neuroscience*, 2003) studied competitors at the World Memory Championships and reported that they did not differ from controls in general cognitive ability or gross brain structure, but did differ in which regions they engaged during encoding, with spatial-navigation areas prominent; the great majority reported using the method of loci. Dresler and colleagues (*Neuron*, 2017) trained novices over several weeks and reported substantial improvement in list recall, partial persistence at a months-later follow-up, and shifts in functional connectivity toward the expert pattern. **In this chapter's vocabulary, both studies are about `eta` and `delta`, not about capacity.**

What is well supported and usually forgotten. Ericsson, Chase and Faloon (*Science*, 1980) trained a single subject over more than 200 hours of practice; his digit span rose from about 7 to about 79. His *letter* span did not move — it remained around six. **The stored order was specific to the essence set it was built on.** In the master table of CANON_INTERNAL §5, `Q` is stock, and stock is indexed by the category set that generated it. There is no general memory to be improved; there are particular `Q`s built against particular `N`s. Any claim that a mnemonic practice raises "memory" as such is a category error the framework catches immediately.

What the framework contributes and the literature does not. The literature reports outcomes. The storage law says *why the outcome has the shape it has*: the gain is a product of three factors, the largest of which is retention, and each is separately attackable. **FVP-5 (identifiability) is the gate such a study would have to pass**, and it is not obvious that it does: two different `(eta, delta)` pairs with the same ratio produce the same steady state, and only the *transient* separates them. The measurement design must therefore observe the approach to steady state and not only the endpoint — which is Chapter 3's method, and why Chapter 3 insists on a zero-input window.

What is not supported by anything here. That the images or places have significance of their own. That a particular palace layout is better than another for reasons other than distinctness and familiarity. That the technique transfers across domains — Ericsson's subject says otherwise.

59.8 One impulse against the same total, computed

This is the calculation the chapter exists to perform. Chapter 2 gave the ceiling argument. Here the comparison is made without any ceiling first, so that the honest baseline is on the record, and the ceiling is added afterwards.

Setup. `delta = 0.06` per week, so `a = 0.94`, `tau_exact = 16.1615` weeks, `t_half = 11.2023` weeks. A fixed total input budget `B = 26.0` units. Three ways to spend it:

- **Impulse:** all 26.0 in week 0, nothing after.
- **Weekly:** 0.5 per week for 52 weeks, nothing after.
- **Four-weekly:** 2.0 every fourth week for 52 weeks, nothing after.

week	impulse	weekly	four-weekly
0	0.0000	0.0000	0.0000
1	26.0000	0.5000	2.0000
2	24.4400	0.9700	1.8800
4	21.5952	1.8271	1.6612
8	16.8604	3.2536	2.9581
13	12.3739	4.6053	5.7325
26	5.5357	6.6655	7.0583
39	2.4765	7.5872	7.3818
52	1.1079	7.9995	7.2731
78	0.2217	1.6010	1.4556
104	0.0444	0.3204	0.2913
156	0.0018	0.0128	0.0117

The crossover is week 24. At week 23 the impulse is still ahead, 6.6648 against 6.3254; at week 24 the cadence is ahead, 6.4458 against 6.2649, and it never trails again while the budget lasts.

And now the fact the folklore gets wrong. Summing `Q_t` over 400 periods:

```
impulse      433.3333
weekly       433.3333
four-weekly  433.3333
B/delta   =  433.3333
```

Identical to every digit. This is not an artefact of the parameters. It is the defining property of a linear time-invariant system: the integral of the response to a given total input is `B/delta` regardless of how the input is distributed in time. **Anyone who tells you that spreading the same total input over more occasions produces more total stored order, in a linear model, has not done the arithmetic.** Chapter 2 said this plainly and this chapter repeats it, because the correct argument is stronger when the incorrect one has been cleared away.

59.9 The steady-state difference — the actual result

The integral is invariant. **The steady state is not, and this is where the difference lives.**

A single impulse of any size whatever, never repeated, has

```
Q_inf = 0
```

because the driven term is zero for all `t > 0` and the transient decays as `a^t`. A sustained rate `m` per period has

```
Q_inf = m / delta
```

Tabulated at `delta = 0.06`:

sustained rate `m` per week	`Q_inf = m/delta`	impulse of any size, unrepeated
0.10	1.6667	0.0000
0.25	4.1667	0.0000
0.50	8.3333	0.0000
0.7278	12.1300	0.0000
1.00	16.6667	0.0000

That column of zeros is the mathematical case for regular practice, and it does not depend on the ceiling, on motivation, on commitment, or on any claim about human nature. A quantity that is not driven decays to zero because `delta > 0`. The only way to hold a nonzero level is to keep driving it. This is why Chapter 1 insisted that the leak is the point: the same `delta` that makes iteration arrive is the `delta` that makes a single arrival not stay.

Now repeat the impulse. Take the same annual total, 26.0 units, delivered forever at interval `P` with burst `A = 26P/52`. The steady-state cycle is exact (Chapter 2, §2.13):

cadence `P`	burst `A`	`a^P`	`Q_max`	`Q_min`	ripple	cycle mean	% of cycle with `Q >= 7`
1	0.5000	0.940000	8.3333	7.8333	0.5000	8.3333	100.0%
2	1.0000	0.883600	8.5911	7.5911	1.0000	8.3333	100.0%
4	2.0000	0.780749	9.1220	7.1220	2.0000	8.3333	100.0%
8	4.0000	0.609569	10.2451	6.2451	4.0000	8.3333	76.9%
13	6.5000	0.447365	11.7618	5.2618	6.5000	8.3333	64.5%
26	13.0000	0.200136	16.2528	3.2528	13.0000	8.3333	52.4%
52	26.0000	0.040054	27.0849	1.0849	26.0000	8.3333	42.1%
104	52.0000	0.001604	52.0836	0.0836	52.0000	8.3333	31.2%

The mean column is constant at 8.3333 in every row. The minimum falls from 7.83 to 0.08. **If a system has any functional threshold at all, the mean is the wrong statistic and the minimum is the right one.** A community whose shared context must stay above some level for its meetings to work, its obligations to hold, or its newcomers to be absorbed does not experience the mean; it experiences the trough, and the trough is the only column in that table that moves.

59.10 The three quantities that are not invariant

Collecting the finite-budget results from §59.8, with thresholds applied:

threshold `theta`	periods with `Q >= theta`, impulse	weekly	four-weekly
3.0	35	60	61

threshold `theta`	periods with `Q >= theta`, impulse	weekly	four-weekly
5.0	27	45	45
7.0	22	25	23
10.0	16	0	0
15.0	9	0	0

Read that table honestly, because it does not say what a partisan of cadence would like it to say. At low thresholds the cadence wins by a wide margin: 60 or 61 periods against 35. At `theta = 7` the three are close. **At `theta = 10` and above the impulse wins outright, and the cadence scores exactly zero, because no cadence funded by that budget ever reaches 10 at all.**

That is the entire formal content of the intensity argument, and it is real. **An impulse buys height that a cadence cannot buy; a cadence buys duration that an impulse cannot buy.** If — and only if — there exists a threshold that must be crossed and that no sustainable rate can reach, the impulse is the correct instrument, and the framework says so without hedging. §59.14 identifies the class of rites that live in exactly that regime and computes what they require.

Three quantities, then, are not invariant to how a fixed budget is spread:

1. **The peak.** `26.0000` for the impulse, `7.9995` for the weekly cadence — a factor of 3.25.
2. **The minimum**, and hence time above any threshold below the cadence's plateau.
3. **The deliverable total itself**, once the ceiling is applied — which is §59.11 and is the largest effect of the three by an order of magnitude.

59.11 The ceiling, applied

Paper D, Corollary 2.2 caps `F(u)` at `ln N`, so per-period input is capped at `eta ln N`. With `eta = 0.35` and `N = 8`:

```
eta ln N = 0.35 * 2.0794 = 0.7278 per week
```

The 26-unit annual total is now attempted at each cadence, and the bursts are clipped:

cadence P	burst wanted	delivered	clipped?	annual delivered	% of weekly
1	0.5000	0.5000	no	26.0000	100.0%
2	1.0000	0.7278	**yes**	18.9229	72.8%
4	2.0000	0.7278	**yes**	9.4615	36.4%
8	4.0000	0.7278	**yes**	4.3668	16.8%
13	6.5000	0.7278	**yes**	2.9112	11.2%
26	13.0000	0.7278	**yes**	1.4556	5.6%
52	26.0000	0.7278	**yes**	0.7278	**2.8%**

The annual rite delivers 2.8% of the weekly practice, from an identical nominal budget. The other 97.2% is not delivered late, or delivered inefficiently. It is not delivered at all: the law caps what one interaction can generate and the excess has nowhere to go.

Design consequence, stated once and load-bearing: an annual event and a weekly practice are not two speeds of the same thing. They are different instruments with different maximum outputs, and the comparison "same total time" is meaningless without the clipping calculation.

59.12 The optimal cadence, and the term that stops the division

Theorem 59.1 says division never hurts. On its own it would recommend infinitely many infinitesimal occasions, which is obviously wrong, and the theorem's honesty requires naming the term that stops it: **the per-occasion fixed cost.**

Let a fixed budget `T = 12` hours be spread over `n` periods with one session of `T/n` hours each. Let non-parallelism rise linearly with session length to saturation at `L* = 8` hours, `u(L) = min(1, L/L*)`. Let `N = 8`, `eta = 0.35`, and let each occasion carry a fixed cost `c_0 = 0.15` regardless of length — convening, travelling, settling, closing.

n	L = T/n (h)	u	f(u)	eta f	delivered total	− n c_0	net
1	12.0000	1.0000	2.0794	0.7278	0.7278	-0.1500	0.5778
2	6.0000	0.7500	1.8326	0.6414	1.2828	-0.3000	0.9828
3	4.0000	0.5000	1.5041	0.5264	1.5793	-0.4500	1.1293
4	3.0000	0.3750	1.2879	0.4507	1.8030	-0.6000	1.2030
5	**2.4000**	**0.3000**	**1.1314**	**0.3960**	**1.9800**	**-0.7500**	**1.2300**
6	2.0000	0.2500	1.0116	0.3541	2.1244	-0.9000	1.2244
8	1.5000	0.1875	0.8383	0.2934	2.3473	-1.2000	1.1473
12	1.0000	0.1250	0.6286	0.2200	2.6402	-1.8000	0.8402
24	0.5000	0.0625	0.3629	0.1270	3.0484	-3.6000	-0.5516
48	0.2500	0.0312	0.1978	0.0692	3.3235	-7.2000	-3.8765

The **delivered total** column is monotone increasing exactly as Theorem 59.1 requires, rising from 0.7278 to 3.3235 — **a factor of 4.57 from the same twelve hours, purely by division.** The **net** column peaks at `n = 5` and falls away on both sides.

Note also the first row. At `n = 1`, the session is 12 hours but u saturates at 8; **four of the twelve hours deliver nothing at all.** The single long occasion wastes a third of the budget before any of the other effects are considered.

The optimum is a function of the per-occasion cost and of nothing else here:

fixed cost `c_0` per occasion	optimal n	session length	net delivered
0.00	(unbounded)	-> 0	-> 3.6275
0.02	24	0.500 h	2.5684
0.05	13	0.923 h	2.0438
0.10	8	1.500 h	1.5473
0.15	5	2.400 h	1.2300
0.25	3	4.000 h	0.8293

fixed cost c_0 per occasion	optimal n	session length	net delivered
0.40	2	6.000 h	0.4828
0.60	1	12.000 h	0.1278
0.80	1	12.000 h	-0.0722

Read the table as a design instrument, not a result. The optimal cadence is not a fact about the practice; it is a fact about how expensive it is to convene. **Halving the cost of convening moves the optimum from five occasions to eight, and doing so is very often cheaper than any intervention on the occasions themselves.** This is the quantitative form of a familiar observation: groups that meet in a place they are already in, at a time they are already free, meet far more often, and the reason is the c_0 column.

At c_0 = 0.80 the net is negative at every n. **Some practices cannot be afforded at any cadence, and the arithmetic says so before anybody's calendar is consulted.**

59.13 Ritual as periodic forcing: what Durkheim and Rappaport are describing

Durkheim's *Elementary Forms of Religious Life* (1912) proposes that periodic assembly generates a state — *effervescence collective* — that the assembled individuals do not generate separately, and that the periodic re-assembly is what keeps the resulting representations alive between times.

In this framework that is two claims, and they map cleanly. The first is that S = F(u) > 0 requires more than one non-parallel participant — Paper A's Theorem 4.3 says surplus is exactly zero when the participating framings are parallel, so an assembly of genuinely different people generates what none of them generates alone. The second is the storage law itself: what the assembly produces is an *input*, and inputs decay unless repeated. **Durkheim's insistence that the feast must recur is the observation that delta > 0.**

Roy Rappaport's *Ritual and Religion in the Making of Humanity* (1999) is more precise and more useful here. Rappaport defines ritual by its form rather than its content — as the performance of sequences of formal acts and utterances that are largely invariant and not encoded by the performers themselves — and separates two streams of information carried by a performance:

- the **canonical stream**, which is the invariant liturgical order itself and is not about the performer at all;
- the **self-referential stream**, which is what this particular performer, at this particular time, transmits by showing up and doing it.

That distinction is a distinction between Q and S. The canonical order is stored order: it persists across performers, changes slowly, and is maintained rather than created each time. The self-referential message is the period's input. Rappaport's central claim — that to perform an order is to accept it, whatever one privately believes — is, formally, a statement that this particular eta is unusually reliable: participation converts into a public commitment record whether or not the participant intends it to.

Rappaport is also explicit that the propositions at the top of a liturgical order — his **Ultimate Sacred Postulates** — are typically unfalsifiable, and that their unfalsifiability is functional rather than concealed. **That is this book's badge rule, arrived at independently and stated in another**

vocabulary. A postulate declared as a postulate does work that a postulate disguised as a finding cannot do, and Rappaport's point is that the declaration is part of the mechanism.

Two experimental results attach here because they concern the input side rather than the storage side. Konvalinka and colleagues (*PNAS*, 2011) recorded heart-rate dynamics during a fire-walking ritual in San Pedro Manrique, Spain, and reported synchronised arousal between performers and related spectators. Wiltermuth and Heath (*Psychological Science*, 2009) reported that synchronous action raised cooperation in subsequent economic tasks. **Both are evidence about S — about what an occasion generates — and neither is evidence about `delta`.** Keeping those two questions apart is most of what this framework contributes to reading the ritual literature.

59.14 The two modes, and what a rare rite must buy

Harvey Whitehouse's *Modes of Religiosity* (2004) distinguishes a **doctrinal** mode — frequent, routinised, low-arousal, transmissible at large scale — from an **imagistic** mode — rare, high-arousal, often painful or frightening, producing intense and enduring episodic memory in small groups. Whitehouse and Lanman (*Current Anthropology*, 2014) develop the second into an account of *identity fusion*.

This is precisely the cadence/intensity axis, arrived at inductively from the ethnographic record, and the framework says something specific about it that the ethnography does not.

A rite at interval `P`, delivering a burst `A`, has steady-state trough `Q_min = A a^P / (1 − a^P)`. Requiring `Q_min >= theta` and solving for the pole:

```
a^P  >=  r / (1 + r),      r = theta / A
```

With the burst at the ceiling, `A = eta ln N = 0.7278`, and a modest functional floor `theta = 0.30`:

interval `P`	in human terms	required pole `a`	required `delta`/week	required `t_half`	in years
1	weekly	0.291884	0.708116	0.56 wk	0.01
4	monthly	0.735026	0.264974	2.25 wk	0.04
13	quarterly	0.909625	0.090375	7.32 wk	0.14
26	half-yearly	0.953743	0.046257	14.64 wk	0.28
52	annual	0.976597	0.023403	29.27 wk	0.56
208	every four years	0.994097	0.005903	117.08 wk	2.25
520	every ten years	0.997635	0.002365	292.71 wk	**5.63**
1560	once a generation	0.999211	0.000789	878.12 wk	**16.89**

A once-a-decade rite requires the trace it leaves to have a half-life of about five and a half years. A once-a-generation rite requires seventeen.

That is the whole explanation of the imagistic mode, and it is arithmetic. **A rare rite is not an inefficient frequent rite. It is a rite operating in a regime where the ordinary `delta` of shared**

practice would leave nothing at all, and it must therefore purchase an extraordinary `delta` by whatever means are available: ordeal, danger, pain, beauty, terror, secrecy, or singularity. The ethnographic observation that rare rituals are overwhelmingly the intense ones is exactly what this table predicts. It is also why they are small: the means of purchasing a five-year half-life do not scale, and Whitehouse's account of why imagistic communities stay small is the same constraint seen from the sociological side.

The requirement is not weakly sensitive to the floor:

floor `theta`	as % of burst	required `delta`/week at `P = 520`	`t_half` in years
0.10	13.7%	0.004056	3.28
0.20	27.5%	0.002947	4.52
0.30	41.2%	0.002365	5.63
0.50	68.7%	0.001726	7.72
0.7278	100.0%	0.001332	10.00

Xygalatas and colleagues (*Psychological Science*, 2013), studying the Thaipusam Kavadi in Mauritius, reported that participants in the higher-ordeal rite subsequently gave more to an anonymous charitable donation than participants in a low-ordeal rite. **The framework offers no opinion on whether that is the right measure, and it is one field result in one population.** What it supplies is a reason to expect the *direction*: the high-ordeal condition is the one that can pay for a low `delta`, and what is measured downstream is a trace that has not yet decayed.

The refusal that governs this whole section, per CANON §5.4: none of this recommends an ordeal. The arithmetic explains why rare intense rites take the form they do. It does not license anyone to build one, and an intervention that removes a person's ability to decline is out of scope for this manual regardless of what its parameters look like.

59.15 The costly-signal break-even

Costly-signalling accounts of ritual (Irons, 2001; Sosis and Alcorta, *Evolutionary Anthropology*, 2003) argue that expensive requirements are hard-to-fake signs of commitment, following Zahavi's handicap logic (1975): a signal is credible in proportion to what it would cost a non-committed sender. Sosis and Bressler (*Cross-Cultural Research*, 2003) examined nineteenth-century American communes and reported that religious communes outlasted secular ones, and that within the religious sample the number of costly requirements was associated with longevity — an association not found in the secular sample. **This is historical and correlational and must be read as such**: no randomisation, one time and place, a survivorship-prone record.

The framework's contribution is not to confirm this but to price it. A costly requirement enters the storage law as an increase in `C`, and can only be worth having if it changes something else. The most plausible candidate is `delta`: a requirement that screens out low-commitment members reduces turnover, and turnover is the dominant term in institutional `delta` (Chapter 23).

From `Q_inf = (eta F − c)/delta`, requiring `Q_inf' >= Q_inf` after adding cost `dc` and moving `delta -> delta'`:

```
(eta F − c − dc) / delta'  >=  (eta F − c) / delta
```

```
dc  <=  (eta F - c) (1 - delta'/delta)
```

The maximum affordable requirement is the current net drive times the fractional reduction in decay it achieves. Worked at `eta = 0.35`, `N = 8`, `u = 0.55` so `F = 1.5790`, `eta F = 0.5526`, `c = 0.15`, `delta = 0.06` per week, giving a baseline net drive of 0.4026 and `Q_inf = 6.7107`:

new `delta`	reduction	max affordable `dc`	as % of net drive
0.055	8.3%	0.0336	8.3%
0.050	16.7%	0.0671	16.7%
0.045	25.0%	0.1007	25.0%
0.040	33.3%	0.1342	33.3%
0.030	50.0%	0.2013	50.0%
0.020	66.7%	0.2684	66.7%
0.010	83.3%	0.3355	83.3%

The percentages in the last two columns are identical, exactly, in every row. That is not a coincidence of the parameters; it falls straight out of the algebra. **A ritual requirement is worth imposing if and only if the fraction of the net drive it consumes is less than the fraction of the decay it removes.** That is a one-line test, it needs two measurements, and it is the sharpest thing this manual has to say about why demanding institutions sometimes outlast easy ones and sometimes merely exhaust their members.

Worked the other way: a requirement costing 0.05 per week must reduce `delta` by at least 12.4% to pay for itself. If `delta` halves, requirements up to 0.2013 per week are profitable and anything above that is not:

added cost `dc`	`Q_inf'` at `delta' = 0.03`	vs baseline 6.7107
0.000	13.4214	better
0.050	11.7548	better
0.100	10.0881	better
0.150	8.4214	better
0.200	6.7548	better (just)
0.283	3.9881	**worse**
0.350	1.7548	**worse**

The failure mode is visible in the table and it is common: a community that has discovered requirements work keeps adding them past the break-even and cannot tell, because the requirements that worked and the requirements that are now destroying it look identical from inside. The only way to distinguish them is to measure `delta` and `C` separately, which is Chapters 3 and 5.

59.16 The multi-rate calendar

Real liturgical calendars are not single-rate. They superpose a fast cycle (daily or weekly), a medium cycle (monthly, or a lunar month), and a slow cycle (annual, or multi-annual). **In cir-**

cuit terms this is a bank of forcings at different frequencies driving one low-pass filter, and Chapter 2's frequency-response table (§2.12) already tells you what happens: the filter passes the slow component and attenuates the fast one, while the ceiling caps whatever arrives in any single period.

The consequence is specific and is almost always missed in practice. Computed at `delta = 0.055` per week, `eta = 0.35`, `N = 8`, ceiling `eta ln N = 0.7278`, cost `0.13` per week, with a weekly session at `u = 0.55` (`0.5526`), a monthly seminar at `u = 0.62` (`0.5863`), and an annual retreat at `u = 0.86` (`0.6821`):

```
  weekly + seminar in the same week  = 0.5526 + 0.5863 = 1.1390
  ceiling                                              = 0.7278
DISCARDED                              = 0.4112   (36.1% of the week's input)

  weekly + retreat in the same week  = 0.5526 + 0.6821 = 1.2347
  DISCARDED                                            = 0.5069    (41.1%)
```

Two good sessions in the same period are not two sessions. Beyond the ceiling they are one. In a three-year simulation the annual retreat placed in a seminar week contributed a steady-state mean of 8.4806; moved to a clean week it contributed 8.5418 — a difference of 0.0612, small in absolute terms and worth having for free, and larger in systems where the two events are closer in size.

Design rule: schedule occasions to avoid collision in the same period before scheduling them for convenience. The calendar is a resource-allocation problem against a per-period cap, and the traditional practice of distributing feasts across the year rather than clustering them is, in this reading, exactly the right allocation whatever its original rationale.

59.17 Open loop and closed loop

One distinction remains, and the formalism makes it more sharply than the ritual literature does.

A liturgical calendar is open-loop forcing. It fires on a date, regardless of the state of the system, and everything in §59.9 applies to it.

Some ritual cycles are closed-loop. Rappaport's earlier ethnography, *Pigs for the Ancestors* (1968), analysed the ritual cycle of the Tsembaga Maring of highland New Guinea as a regulator: the great festival was not calendrical but was triggered when the pig herd had grown to the point of straining the households that tended it, and the cycle thereby governed pig population, warfare and land use together. Rappaport described it in explicitly cybernetic terms.

A state-triggered rite is a controller, and controllers have failure modes that calendars do not. Chapter 2's §2.16 gives them: a loop acting on a stale measurement overshoots — 30.2% at a lag of two periods, 75.1% at four — and at sufficient gain and lag it diverges outright.

The two designs trade different risks and both risks are computable. A calendar cannot respond to conditions and will sometimes fire when there is nothing to fire about. A trigger responds but inherits every pathology of delayed feedback. Systems that run both — a fixed calendar for the base rate, a triggered rite for exceptional states — are hedging exactly as the control theory recommends, and they are extremely common.

59.18 The spacing effect, as the empirical counterpart

The claim that distributed practice outperforms massed practice for the same total study time is one of the oldest and best-replicated findings in experimental psychology. Ebbinghaus reported it in *Über das Gedächtnis* (1885), and the quantitative synthesis by Cepeda, Pashler, Vul, Wixted and Rohrer (*Psychological Bulletin*, 2006) reviews the verbal-recall literature and finds the advantage robust across a wide range of conditions. Cepeda and colleagues (*Psychological Science*, 2008) further reported an interior optimum: the best gap between study episodes grows with the interval until test, at a ratio commonly reported in the range of roughly ten to twenty per cent.

Two honest statements, in order.

First, the spacing effect and Theorem 59.1 are not the same claim and must not be presented as mutual confirmation. The spacing literature is about human memory encoding; Theorem 59.1 is a statement about concave functions. **What they share is a structural prediction — that division of a fixed budget into more occasions helps, up to a point set by a per-occasion cost — and structural agreement between an empirical regularity and a mathematical result is interesting without being evidence for either.**

Second, the finding that the optimal gap is a roughly fixed *fraction* of the retention interval says that the right cadence scales with the system's own time constant rather than with the calendar — which is the content of Chapter 2's specification S2.7, `P <= tau/2`. **Both put the practice interval at a fixed ratio of the system's own timescale, and both put that ratio well below one. That agreement is structural and this manual will not claim more from it than that.**

Roediger and Karpicke (*Psychological Science*, 2006) add a second empirical handle: retrieval practice — being tested — produces better long-term retention than additional study, and the advantage grows with the retention interval. **Formally this is an `eta` result, not a `delta` result: the same occasion converts more of its content into durable structure when it requires retrieval rather than exposure.** The ritual counterpart is recitation rather than reading, catechism rather than sermon, examination rather than attendance — and the tradition's near-universal preference for the first of each pair is, on this reading, a preference for the higher-`eta` form of the same occasion.

The Specification

S59.1 — Declare the retention medium. Every stored-order design shall state what the medium of retention is (memory, place, object, document, personnel, habit) and record its `delta` with its period, per S2.1. Media with different `delta` shall be modelled as separate stocks, not averaged.

S59.2 — Price a palace by doublings. Palace capacity shall be recorded as `ln L` in the number of loci. Any proposal to expand a palace shall state the number of doublings purchased and the construction and maintenance cost of each; each doubling delivers exactly `ln 2 = 0.6931` nats.

S59.3 — Record the three-factor decomposition. Any claim that an encoding technique improves storage shall report the claimed movement in `u`, in `eta`, and in `delta` separately, and shall show their product against the claimed `Q_inf` ratio. A claim that reports only the outcome is not auditable and shall be recorded as unverified.

S59.4 — Separate `u_cross` from `u_div`. Where both between-locus and within-locus distinctness are being engineered, both shall be estimated and Theorem 4.1's bound `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)` shall be shown. The additive figure shall never be used as the estimate.

S59.5 — State the essence set that `Q` is built against. Stored order is indexed by its category set. No claim of general improvement in a faculty shall be made; the claim shall name the `N` it was built on (Ericsson, Chase and Faloon, 1980).

S59.6 — Compute the integral before arguing about cadence. Any cadence proposal shall show that the integrated effect of the alternatives is equal in the unclipped linear model, and shall then identify which of the three non-invariant quantities — peak, minimum, or deliverable total — the proposal is actually trading.

S59.7 — Report `Q_min`, never only the mean. Per S2.8. Additionally, the functional threshold `theta` below which the system is observed to misbehave shall be stated, and the percentage of the cycle spent above it shall be computed.

S59.8 — Clip every burst at `eta ln N` before comparing totals. Per S2.6. The annual delivered total shall be computed after clipping, and the discarded fraction shall be reported explicitly.

S59.9 — Size occasions by the fixed-cost optimum. The per-occasion fixed cost `c_0` shall be measured, the optimum n computed against it, and the sensitivity of the optimum to `c_0` shown. Any proposal to increase cadence shall first consider reducing `c_0`.

S59.10 — Justify a rare rite by its required `delta`. Any occasion at interval `P > tau` shall compute the retention half-life its trace must have to hold the stated floor, using `a^P >= theta/(theta + A)`, and shall state by what means that retention is expected to be achieved. If no such means can be named, the occasion shall be reclassified as ceremonial and shall not be counted as an input to `Q`.

S59.11 — Apply the costly-requirement break-even. No requirement shall be added to a practice without recording its cost `dc`, the current net drive `eta F − c`, and the reduction in `delta` it is expected to produce. The requirement is admissible only if `dc <= (eta F − c)(1 − delta'/delta)`, and the expected `delta'` shall be measured afterwards, not assumed.

S59.12 — Deconflict the calendar. Scheduled occasions shall be checked for collision within a single period. Where two occasions fall in one period and their combined input exceeds `eta ln N`, the discarded amount shall be computed and one occasion shall be moved unless a reason for the collision is recorded.

S59.13 — Declare open-loop or closed-loop. Every recurring occasion shall be labelled calendrical (open-loop) or state-triggered (closed-loop). Closed-loop occasions shall additionally record the measurement lag L and correction gain k and satisfy S2.10.

S59.14 — Prefer the retrieval form of an occasion. Where an occasion can be run as exposure or as retrieval at equal cost, the retrieval form shall be specified, and the `eta` difference shall be recorded as an estimate to be tested rather than as an established value.

The Failure Modes

THE HEROIC INTENSIVE — a large single occasion substituted for a practice, justified on total input.

- *Parameter damaged:* delivered `S`, clipped at `eta ln N`, while `C` continues unabated.
- *Diagnostic signature:* a strong peak, a crossover against the cadence alternative at roughly `1.5 tau`, and a post-mortem that concludes "we failed to sustain it." In §59.8's parameters the crossover is week 24 against `tau = 16.16`.
- *Measurement:* compute the clipped annual delivered total at each cadence (§59.11 table). Compute `Q_inf = 0` for the unrepeated impulse and `m/delta` for the alternative.
- *Fix:* redistribute to the cadence whose per-period input sits at or just under the ceiling. Retain the intensive only if a threshold has been identified that no sustainable rate reaches — and if so, say which threshold, in the design record.

THE MEAN-ONLY CADENCE ARGUMENT — a sparse schedule defended on the grounds that the average stored order is unchanged.

- *Parameter damaged:* none. The mean really is unchanged. The minimum is not.
- *Diagnostic signature:* the schedule's advocates cite the total or the average; its participants describe "getting back up to speed" each time. In §59.9's table the mean is 8.3333 in every row while `Q_min` falls from 7.8333 to 0.0836.
- *Measurement:* compute `Q_min = A a^P/(1−a^P)` and the percentage of the cycle above the stated functional threshold.
- *Fix:* shorten `P`, or state the threshold and demonstrate that the trough clears it. **Averages do not fail systems; minima do.**

THE UNPAID RARE RITE — an occasion at an interval far longer than `tau`, treated as an input to stored order, whose trace cannot survive to the next occurrence.

- *Parameter damaged:* effectively `eta`, since the converted order is gone before it can compound.
- *Diagnostic signature:* the occasion is memorable and well-regarded, and the system's behaviour between occurrences is indistinguishable from a system that never held it. In §59.16's simulation an annual-only schedule holds `Q > 0` for four weeks of fifty-two.
- *Measurement:* compute the required `delta` from `a^P >= theta/(theta+A)` and compare against the measured `delta`. In the worked example the shortfall factor is 2.5.
- *Fix:* either supply the retention the interval demands — an artefact, a record, a durable external store, per §59.6's `delta` column — or shorten the interval, or reclassify the occasion honestly as ceremony rather than as input.

THE COLLIDING CALENDAR — two substantial occasions scheduled into the same period, so that the second is largely discarded at the ceiling.

- *Parameter damaged:* delivered `S` in the colliding period.
- *Diagnostic signature:* a heavy week followed by a light one; participants report that the second event "did not land." The discarded fraction is computable and is often above a

third.

- *Measurement:* sum the intended `eta F` of all occasions falling in one period and compare against `eta ln N`. §59.16 computes 36.1% and 41.1% discard for two realistic collisions.
- *Fix:* move one occasion to a clean period. This costs nothing and is the highest return-per-effort intervention in this chapter.

THE REQUIREMENT RATCHET — costly requirements added past the break-even, because earlier requirements demonstrably helped.

- *Parameter damaged:* `C`, without a compensating fall in `delta`.
- *Diagnostic signature:* membership requirements have grown monotonically for years; retention among those who remain is excellent; recruitment and total membership are falling. The institution interprets falling numbers as evidence that standards are working.
- *Measurement:* apply `dc <= (eta F − c)(1 − delta'/delta)` to each requirement in force. Measure `delta` before and after each addition rather than assuming the direction.
- *Fix:* remove requirements from the bottom of the value ordering until the inequality holds, and measure. Note the asymmetry Chapter 23 records: **an institution that will not choose what it sheds does not thereby keep it.**

THE PALACE THAT WAS NEVER MAINTAINED — a memory system built once, whose loci themselves have decayed, so that retrieval fails for reasons the operator attributes to the contents.

- *Parameter damaged:* `delta` of the index rather than of the material.
- *Diagnostic signature:* recall fails specifically at the transitions between loci; the operator can produce isolated items but not the ordered walk. Diagnostic separation from ordinary forgetting is that prompting with the *place* restores the item.
- *Measurement:* test the empty palace — walk it with no contents and count loci recovered in order. Compare against the count when it was built.
- *Fix:* rebuild the index before re-loading the contents. Loading new material onto a decayed index raises `C` and lowers `eta` simultaneously and is the worst available move.

THE UNDECLARED ESSENCE SET — improvement measured on one category set, claimed for a faculty.

- *Parameter damaged:* none; the claim is simply not about the quantity measured.
- *Diagnostic signature:* a practice that produced a large, real gain in one domain produces nothing in a second, and the operator concludes the second domain is harder rather than that the stock is domain-indexed.
- *Measurement:* run the same protocol on a disjoint essence set. Ericsson, Chase and Faloon's subject reached a digit span near 79 while his letter span stayed near six.
- *Fix:* declare the `N` in every claim. Build a second stock deliberately if a second domain is wanted; there is no transfer to appeal to.

Worked Example

A chartered study circle: the palace and the calendar, priced together

The system. Six people from six disciplines convene as a chartered study circle around a canonical corpus of 48 items. There are two coupled subsystems with different natural periods, and they must be modelled separately per S59.1: each member's individual hold on the corpus (period: one day) and the circle's own stored order (period: one week).

The question on the table is whether to (a) adopt loci encoding for the corpus, (b) replace the current monthly three-hour seminar with a weekly ninety-minute session, and (c) keep or move the annual retreat.

All parameter values below are design values, badged as such, obtained by the protocols of Chapters 3, 4 and 9 in a real engagement and stipulated here.

Step 1 — Subsystem M: one member's hold on the corpus. Period: one day.

```
N = 48 items          ln N = 3.8712 nats     <- hard per-occasion ceiling
```

	u	N_eff	F	% of ceiling	eta	delta/day	eta5	Q_inf	DC gain	t_half	t_95
before: reread only	0.20	10.400	2.3418	60.5%	0.18	0.22	0.4215	1.9160	0.818	2.79 d	12.06 d
after: 24-locus palace, two images each	0.88	42.360	3.7462	96.8%	0.36	0.05	1.3486	26.972	77.200	13.51 d	58.40 d

```
Q_inf ratio = 14.0774  =  F 1.5997  x  eta 2.0000  x  1/delta 4.4000
```

Per S59.3 the three factors are reported separately, and per S59.4 the split is shown. Estimating `u_cross = 0.70` (distinct places) and `u_div = 0.18` (striking images):

```
f(u_cross) = 3.5234     f(u_div) = 2.2471     f(u) = 3.7462
bound:  3.5234  <=  3.7462  <=  5.7705
```

Place alone accounts for 94.1% of the combined surplus. The images are worth having and they are worth 0.2228 nats, not the 2.2471 that the additive reading would suggest. **Per S59.4 the additive figure is not used, and a design that spent equally on images and places would be spending most of its effort on the smaller term.**

Per S59.5, the claim on record is: *ordered recall of this 48-item corpus improves.* Not memory.

Step 2 — Subsystem G: the circle. Period: one week.

The circle is six people, but the category set that matters is the set of *domains* they bring, and with two recurring outside guests the Chapter 7 enumeration returns eight:

```
N = 8 domains                          ln N     = 2.0794
u* = 0.55 (Chapter 9 audit)            N_eff    = 1 + 7(0.55) = 4.8500
F = ln 4.85                                = 1.5790     (75.9% of ceiling)
eta = 0.35, delta = 0.055/wk, C = 0.13/wk
ceiling eta ln N                                    = 0.7278/wk
```

Rather than assume the session geometry, the cadence question is put to Theorem 59.1 directly. The circle has a genuine interaction budget of **three hours per fortnight** and a per-occasion cost of `c_0 = 0.045` (convening, travel, settling). Non-parallelism saturates at `L* = 2.5` hours.

schedule	occasions per fortnight	length	u	f(u)	eta f(u)	over ceiling?	delivered per fortnight	vs one long session
one 3.0 h session	1	3.000 h	1.0000	2.0794	0.7278	at the ceiling	0.7278	1.0000
two 1.5 h sessions	2	1.500 h	0.6000	1.6487	0.5770	no	1.1541	1.5857
three 1.0 h sessions	3	1.000 h	0.4000	1.3350	0.4673	no	1.4018	1.9260
four 45 min sessions	4	0.750 h	0.3000	1.1314	0.3960	no	1.5840	2.1764
six 30 min sessions	6	0.500 h	0.2000	0.8755	0.3064	no	1.8385	2.5261

Note that the single three-hour session reaches `u = 1` and therefore sits **exactly** at the ceiling, `eta ln N = 0.7278`: it delivers the maximum a single period can carry, and the half-hour of that session beyond `L* = 2.5` delivers nothing at all. Two fortnightly sessions deliver 1.5857 times as much; six deliver 2.5261 times as much. **The delivered column is monotone in n, exactly as Theorem 59.1 requires.**

The circle adopts the weekly ninety-minute form — `n = 2` per fortnight — on the grounds that the higher-n options push per-occasion cost above what six busy people will actually pay, and because the ninety-minute session leaves visible headroom under the ceiling for a guest or a longer thread. Session geometry at 1.5 hours would give `u = 0.60` and `eta f = 0.5770`; the conservative figure carried forward is the *audited* `u* = 0.55` from the Chapter 9 disagreement survey, since that is what the circle actually achieves across a year rather than in a well-run session:

```
weekly input   eta F = 0.35 * 1.5790          = 0.5526/wk
net drive = 0.5526 - 0.13 = 0.4226/wk
Q_inf     = 0.4226 / 0.055 = 7.6844
pole      = 0.945    tau = 17.677 wk    t_half = 12.253 wk
t_95      = 52.96 wk    t_99 = 81.41 wk
h at Q_inf = 0.7278 - 0.055(7.6844) - 0.13 = 0.1752/wk       (r_max = h)
h = 0 at Q = 10.8692
```

Per S2.4 the first honest review is week 53. A week-13 reading of 4.4040 is 51.6% of steady state and is not a finding about the change.

Step 3 — The calendar, deconflicted. Per S59.12.

Three occasions are proposed: the weekly session (`u = 0.55`, `eta F = 0.5526`), a monthly seminar with an outside guest (`u = 0.62`, `eta F = 0.5863`), and an annual retreat (`u = 0.86`, `eta F = 0.6821`). The per-period ceiling is 0.7278.

```
weekly + seminar  = 1.1390    ceiling 0.7278    DISCARDED 0.4112   (36.1%)
weekly + retreat  = 1.2347    ceiling 0.7278    DISCARDED 0.5069   (41.1%)
```

The retreat is initially scheduled for the last week of the year, which is also a seminar week. Three years simulated, with cost paid every week:

schedule	steady mean Q	min	max	ripple
weekly only	7.6844	7.6844	7.6844	0.0%
weekly + monthly seminar + retreat, **retreat colliding with a seminar**	8.4806	—	—	—
weekly + monthly seminar + retreat, **retreat in a clean week**	8.5418	8.4259	8.7145	3.38%

Trajectory of the built calendar:

week	1	4	13	26	53	105	157	209
Q_t	0.4226	1.7313	4.4040	6.6679	8.1134	8.5193	8.5408	8.5419

Moving the retreat off the seminar week costs nothing and returns 0.0612 of steady-state stock. The full calendar returns +0.8574 over weekly-only, an 11.16% gain, and headroom at the steady mean is `0.7278 - 0.055(8.5418) - 0.13 = 0.1280`, still positive. Headroom reaches zero at `Q = 10.8692`, comfortably above the operating point.

Step 4 — Is the annual retreat an input at all? Per S59.10.

Standing alone, at `P = 52`, `A = 0.6821`, and a functional floor `theta = 0.25`:

```
required a^P  >=  theta/(theta + A) = 0.25/0.9321 = 0.2682
required a    >=  0.977559        required delta <= 0.022441/wk
required t_half >= 30.54 weeks = 0.59 years
```

The circle's measured `delta` is 0.055 per week, a half-life of 0.24 years. **The shortfall factor is 2.5.** Formally: *the retreat cannot, by itself, hold a floor of 0.25 for a year at this circle's retention.* Simulated alone, an annual-only schedule holds `Q > 0` for four weeks in fifty-two.

Per S59.10 the retreat is therefore not recorded as an input to `Q` on its own account. It is recorded as (a) a `+0.0612` contribution on top of the weekly base, and (b) a candidate `delta` intervention — the retreat is where the circle's charter is re-read and its obligations restated, which is a retention measure, not an input measure, and it is to be measured as one.

Step 5 — The proposed attendance requirement. Per S59.11.

A member proposes a hard attendance requirement: miss four consecutive weekly sessions and lapse. Estimated additional cost `dc = 0.05` per week in administration, chasing, and the friction of enforcement.

```
net drive            = eta F − C = 0.5526 − 0.13 = 0.4226/wk
required reduction   = dc / net drive = 0.05/0.4226 = 11.83%
required delta'      <= 0.055 ∗ (1 − 0.1183) = 0.048493/wk
```

The requirement pays for itself if and only if it reduces weekly decay by at least 11.8% — in membership terms, if it raises the half-life of continuous participation from 12.25 weeks to at least 13.90. That is a specific, measurable, falsifiable claim, and it is the claim the circle records before adopting the rule, not after.

Per S59.11 the outcome is to be measured. Per the requirement-ratchet failure mode, this is the *first* requirement and the break-even is generous; the test will get harder for the second and harder still for the third, and the circle records that fact in the same minute.

Step 6 — What would falsify all of this. Per Chapter 25 and FVP-6.

Fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on weeks 1-53 with `0 <= a <= 1`, `b >= 0`, freeze the coefficients, and report on weeks 54-106 against four baselines: persistence, direct-input, matched autoregressive, and shuffled-input control. Predictions on record:

1. The fitted a lies in `[0,1]` and its interval contains `1 − delta = 0.9450`. If not, the first-order form is wrong for this circle and every number above is void.
2. The fitted b exceeds its pre-intervention value. If not, the cadence change did not raise delivered input, whatever the participants report.
3. The frozen model beats the **shuffled-input control** on held-out data. **If it does not, then the timing of the circle's occasions carries no information beyond their volume, the entire cadence argument does not apply to this system, and Steps 2 through 5 should be withdrawn.**

4. On Subsystem M, the empty-palace test (per the palace failure mode) recovers at least 20 of 24 loci in order at 90 days. If the index itself decays faster than the contents, the `delta` attribution in Step 1 is misassigned.
5. The attendance requirement reduces measured membership `delta` by at least 11.8% within two `tau`. If it does not, it is removed — not retained on the grounds that standards matter.

Prediction 3 is the decisive one, exactly as in Chapter 2. Everything this chapter claims about cadence is a claim that *when* input arrives matters and not only how much. The shuffled-input control is the test that can kill it, and it is the control most often omitted.

What this chapter has established

The method of loci is the storage law with space as the retention medium. Its advantage factorises into three separately measurable terms — surplus, conversion, and retention — and in the worked case retention supplies more than the other two combined, which is the opposite of the usual emphasis. Palace capacity is `ln L`, so **every doubling of the palace buys exactly 0.693147 nats and no arrangement of the loci changes that.** The classical rules of the *ad Herennium* map one-to-one onto the parameters of the law with no rule left over and no parameter left untouched.

Ritual is periodic forcing of a low-pass filter, and the case for cadence over intensity has two independent halves. **The first is Chapter 2's ceiling: surplus saturates at `ln N` per period while cost does not, so an annual rite delivers 2.8% of what a weekly practice delivers from the same nominal budget.** The second is Theorem 59.1: because `f(u)` is strictly concave through the origin, dividing a fixed interaction budget into more occasions never delivers less and in the worked case delivers 4.57 times as much — a result that follows from Paper A's Theorem 3.3 and needs no ceiling at all. What stops the division is the per-occasion fixed cost, and the optimum is computable from it: **lowering the cost of convening moves the optimum more cheaply than any intervention on the occasions themselves.**

The honest comparison between one impulse and the same total delivered periodically was computed and it is not one-sided. Integrated effect is identical to every digit — 433.3333 in all three regimes — because the system is linear and time-invariant. Peak favours the impulse by a factor of 3.25. Minimum, time above a low threshold, and clipped deliverable total favour the cadence, the last by a factor of 36. **And the steady state settles it: a single impulse, however large, has a steady state of exactly zero, while a rate `m` has `m/delta`.**

The rare, intense rite is not thereby refuted. It occupies the regime above the cadence's plateau, where a threshold exists that no sustainable rate reaches — and the price of that regime is computable. **A once-a-decade rite requires its trace to have a half-life of about five and a half years, and a once-a-generation rite about seventeen.** That is why such rites are extraordinary, why they are small, and why they take the forms they do. It is also the boundary of what this manual will say about them: the arithmetic explains the form and does not recommend it.

Chapter 60 closes the volume. It states the thesis that has been implicit since Chapter 57 — that the old notation and this one describe the same relationships — and it does so under the strictest constraint in the book: **claim no supernatural efficacy, no physical effect of symbol, and no result that does not follow from the formalism.** The argument is stronger for the restriction, because what remains after the restriction is applied turns out to be sufficient.

Chapter 60 . This Is How It Is Done

Fifty-nine chapters ago this manual opened with a difference equation and the claim that its leak was not a defect. Everything since has been that equation at a different scope, with different parameters, under a different name.

This chapter states what the whole assembly amounts to, and it states it under the strictest constraint in the book. **CANON §5 governs absolutely here: no supernatural efficacy is claimed, no physical effect of symbol is claimed, and no result appears that is not derivable from the formalism.** The chapter's force comes entirely from that restriction. A synthesis that permitted itself one unearned step would be worth nothing, because the interesting question is not whether the old notation *can* be dressed in new symbols — anything can — but whether the modern notation is **sufficient**: whether, having described the mechanism, there is a residue that still requires something else.

The finding, stated plainly and then defended for the rest of the chapter, is that there is not — at least not in the audit performed here, over the operations this manual has examined. **That is a claim about an audit and not a proof of completeness, and it is badged as such wherever it appears.**

This is also the end of Volume Two and not the end of the work. Volume Two turned the instrument inward and found the same instrument. **Volume Three turns it downward, onto the biological layer, and that is a genuinely different problem** — the closing section, *The Next Instrument*, says why.

The Law

60.1 The thesis

THESIS (the author's, stated as a postulate and badged as one). If *magic* is understood as

1. the deliberate engineering of the perceptual lens — the ratio `D/P` and the moment `M = (D/P)E`;
2. the storage of order through symbol, space and cadence — the parameters `eta`, `delta`, and the stock `Q`;
3. the composition of non-parallel essences — the geometry `u` and the surplus `F = ln(1 + (N−1)u)`;

then this manual has described its mechanism in modern notation.

That is the whole claim. Note carefully what it contains and what it does not.

It contains a definition. The word *magic* is being given a stipulative reading — three engineering operations — and the thesis says that under that reading, the mechanism is the one this manual has specified. **A reader who rejects the definition is not thereby contradicting the thesis; they are declining it, which is a different act, and the manual has no argument against them.**

It does not contain an empirical claim about the historical practice. Whether the practitioners of any tradition understood themselves to be doing these three things is a question for historians, and this manual does not answer it.

It does not contain an efficacy claim of any kind. The thesis says the mechanism *is described*. It does not say the described mechanism produces effects beyond those the storage law and the surplus law produce, because they are the same mechanism. **There is nothing left over for an effect to come from.**

60.2 What "described its mechanism" means

A mechanism is described when three things hold, and all three are checkable:

```
(i)  every operation in the practice maps to a named parameter of the formalism
(ii)  the direction and, where measurable, the magnitude of each operation's effect
        follows from the law rather than being stipulated alongside it
(iii) the failures of the practice are the failures the law predicts, and appear at
        the parameter values where the law says they will appear
```

§60.18 performs (i) as an explicit audit over eleven classical operations. §60.6 through §60.16 perform (ii) part by part. The Failure Modes sections of all sixty chapters are (iii), and Chapter 27's catalogue is its index.

What would refute the thesis is a fourth thing: an operation that does work in the practice, is agreed to do work, and maps to no parameter. The manual has looked for one and §60.19 records what was found. **That is the honest form of the claim and it is falsifiable in exactly one direction: someone may exhibit the missing operation.**

60.3 The ledger

The single most useful page in this chapter is an inventory of everything the manual has asserted, sorted by status. **A synthesis that does not publish this table is asking to be believed rather than checked.**

PROVED — mathematics, true regardless of whether anything in the world obeys it:

Result	Source
`f(u) = ln(1+(N−1)u)` uniquely characterised by S1-S4	Paper A, Thm 2.1
Strict monotonicity, positivity, strict concavity of `f`	Paper A, 3.1-3.3
Subadditivity on the parameter u (**not** on geometric triples)	Paper A, 3.4
Sharp Lipschitz constant `N−1`	Paper A, 4.2
`F = 0` iff parallel; `F = ln N` iff orthogonal	Paper A, 4.3
Two-source decomposition and its bracket	Paper A, Thm 4.1
`0 <= F(u) <= ln N` — the hard ceiling	Paper D, Cor 2.2
Exact solution of the storage recursion	Paper E, Thm 3.1
One-step growth condition `eta S_t > delta Q_t + C_t`	Paper E, Cor 3.2
Uniform bound `Q_t <= max{Q_0, eta ln N/delta}`	Paper E, Thm 3.3

Result	Source
Sustainability ceiling `C_t <= eta ln N − delta Q_t`	Paper E, Cor 3.4
Steady state `Q_inf = (eta F(u*) − c*)/delta`	Paper E, Thm 3.5
Max one-step increment `r_max = eta ln N − delta Q_t − C_t`	Paper E, Cor 4.2
The recursion is a first-order low-pass filter with pole `1−delta`	Chapter 2
Contraction implies a unique fixed point reached from any start	Banach
`f(n) = k n^2` is the exact collapse of the published essence recipe	Chapter 29
Essence intervals are ratios of squares; `k` cancels	Chapter 29
An essence set `1..N` has exactly `ceil(N/2)` pitch classes	Chapter 29
One nat of surplus is exactly `2/ln 2 = 2.885390` octaves	Chapter 32
A star polygon `{N/k}` has exactly `gcd(N,k)` closed strokes of length `N/gcd(N,k)`	Chapter 58
Dividing a fixed budget over more occasions never delivers less, for concave `g`, `g(0)=0`	Chapter 59, Thm 59.1
Costly-requirement break-even `dc <= (eta F − c)(1 − delta'/delta)`	Chapter 59
Expected surplus of a uniform random correspondence `= (N ln N − N + 1)/(N−1)`	§60.18
Expected colliding pairs under a hash `>= C(M,2)/R`, minimised by uniformity	§60.20

POSTULATED — badged every time, chosen rather than derived:

Postulate	Where
S3, the affine effective count — the *design* axiom of the surplus law	Paper A
That a room, team or institution is described by the storage recursion at all	CANON §1.4
Hyperrecursive symmetry: the same form at every scale	CANON §1.4
That essence types have a frequency; that the law is square in the index	Chapter 29
That `k = (9/8)^2/phi^2` rather than any other constant	Chapter 29
`M = (D/P)E` as the form of the moment	Chapter 34

Postulate	Where
The internal/external correspondence as one subject at two scopes	Chapter 50
The correspondence table read as a nonlinear computing map — **a metaphor**	Chapter 57
That the method of loci raises u and `eta` and lowers `delta`	Chapter 59
The thesis of §60.1	this chapter

REQUIRES MEASUREMENT — asserted nowhere, obtainable by protocol:

Every `delta`, every `eta`, every u, every `C`, every `N`, every `D` and `P`, in every system, for every period. **The manual supplies no value of any of these for any real system.** It supplies the protocols: Chapter 3 for `delta`, Chapter 4 for `eta`, Chapter 9 for u, Chapter 5 for `C`, Chapter 7 for `N`, Chapters 35 and 53 for `D` and `P`, Chapter 24 for all of them at once.

Counted: 24 proved results, 10 postulates, and a parameter set that is entirely empirical. The ratio is the point. **A framework whose postulates outnumber its theorems is a worldview; one whose theorems outnumber its postulates by more than two to one, with every postulate labelled, is an engineering discipline with a declared foundation.**

60.4 The refusals are load-bearing

CANON §5 lists four refusals, restated in the front matter and honoured in every chapter. They are not modesty. **They are the conditions under which the thesis of §60.1 means anything at all.**

1. **No claim that the mathematics proves anything about physics, thermodynamics, or consciousness.** Paper E proves no negentropy law; Corollary 2.2 there gives `F(u) + N_can(u) <= ln N`, which is a bound and not a mechanism.
2. **No postulate without its badge.** §60.3's second table is the badge, collected.
3. **No number that was not computed or cited.** Every figure in this volume was computed; where a parameter is illustrative it says so in the same sentence.
4. **No intervention that removes a person's ability to decline.** This is the refusal that governs Chapter 59's ordeal arithmetic, Chapter 19's contagion engineering, and every use of the word *design* applied to other people.

If any of the four were suspended, the thesis would collapse into an unfalsifiable analogy. With all four in force, it is a specific, checkable, and modest claim: *this is the mechanism; there is nothing else in it; here is the arithmetic.*

The Engineering

60.5 The arc

Three movements, one instrument.

Book One — the operator learns to read the gradient. *Creature Comforts and Their Hidden Power* (2026) is twelve nodes and a story. Its instrument is attention: notice the chair the dream

is sitting in; inventory what your comforts cost and what your discomforts quietly pay; trace a sightline forward without switching assumptions halfway. Its subject is one person, reading.

Volume One — the engineer builds the circumstance. Chapters 1-31 take the same relations and make them constructible: the circuit and its four parameters, the geometry of difference, the room and its acoustics and its thresholds, the sociology of groups and institutions, the tradecraft of surveying and preregistering and intervening, and finally the essence law that gives a relationship a name that can be heard. Its subject is the world outside one person, being built.

Volume Two — the instrument turns on its user. Chapters 32-60 take the same six parameters and apply them at internal scope, and the discovery is not that psychology *resembles* sociology but that **they are one set of relations measured at two scopes** (CANON_INTERNAL §0, badged as thesis). The perceptual lens is `D/P`. Character is `Q`. Forgetting is `delta`. Personal capacity for the new is `h`. The many windows of the one tower are `F_int` over one observer's own framings, which is Paper A applied within a single mind — a derivation, not an analogy.

The movement between the three is not a change of subject. It is a change of who is holding the instrument and what they are pointing it at.

```
BOOK ONE      one person, reading            instrument: attention
              "notice the chair the dream is sitting in"
                        |
                        |   the same relations, made constructible
                        v
VOLUME ONE    the world outside, being built  instrument: specification
             "here is how to build the room, the table, the institution"
                        |
                        |   the same instrument, pointed back
                        v
VOLUME TWO    the builder, being measured     instrument: the formulary
             "the thing you used to build the room is the thing you are"
```

The third movement is the one that closes the loop, and it closes it in a specific and slightly uncomfortable way. An engineer who has spent Volume One measuring other people's `u`, budgeting other people's `h`, and designing cadences for other people's calendars arrives in Volume Two to find the same instruments turned around: their own `N` is the count of things they can actually distinguish; their own `delta` is how fast they lose what they learn; their own `D/P` decides whether they are seeing a system or their model of one. **The manual's diagnostics do not exempt the person holding them, and Chapter 48's over-framed institution and Chapter 35's over-framed observer are the same fault at two scopes.**

There is a further consequence worth stating because it governs how the three movements should be read together. **Book One's instrument was never superseded.** Attention is what produces the survey data of Chapter 24, the disagreement audit of Chapter 9, and the magnitude estimates of Chapter 53. A reader who takes Volume One's specifications without Book One's reading practice has a set of formulas and no measurements to put in them, and the formulas will accept any numbers offered. **The sequence is not a ladder where the earlier rungs are discarded. It is three instruments that must all be working at once, and the manual's most common misuse is running the arithmetic on estimates that nobody actually looked at.**

Two correspondences from Book One deserve restating because they are exact, not deco-

rative:

- **Node 09 RECURSE is the transient `(1-delta)^t Q_0`** — the initial state bleeding off while the driven term builds.
- **Node 10 FIXEDPOINT is Paper E, Theorem 3.5.** The room that was always furnished is the steady state of a contracting map.

And that is why arrival feels the way it does. The contraction ratio is `(1-delta)^t`:

delta	t=1	t=5	t=10	t=20	t=50	t=100
0.020	0.9800	0.9039	0.8171	0.6676	0.3642	0.1326
0.055	0.9450	0.7536	0.5680	0.3226	0.0591	0.0035
0.100	0.9000	0.5905	0.3487	0.1216	0.0052	0.0000
0.200	0.8000	0.3277	0.1074	0.0115	0.0000	0.0000

A contraction does not arrive somewhere new. It arrives where it was always going to arrive, from wherever it started. Recognition rather than victory is the correct phenomenology of a fixed point, and Banach's theorem is why.

60.6 The same equation, six scopes

The architectural claim of this manual is that only `(N, u, eta, delta, C)` and the period change between a desk and a city. That claim is about **model form**, and it is honest because it is checkable: if a scope required a different equation, the claim would fail there.

Computed across seven scopes, with illustrative parameters badged as such:

scope	period	N	u*	eta	delta	C	ln N	N_eff	F	eta ln N	Q_inf	tau	t_95	h
a desk	day	4	0.45	0.30	0.100	0.02	1.3863	2.350	0.8544	0.4159	2.3632	9.49	28.43	0.1596
one mind (a practice)	day	12	0.55	0.30	0.050	0.06	2.4849	7.050	1.9530	0.7455	10.5182	19.50	58.40	0.1596
a dyad	week	6	0.62	0.40	0.040	0.08	1.7918	4.100	1.4110	0.7167	12.1092	24.50	73.39	0.1523
a room / a team	week	8	0.55	0.35	0.055	0.13	2.0794	4.850	1.5790	0.7278	7.6844	17.68	52.96	0.1752
a department	month	14	0.48	0.30	0.050	0.22	2.6391	7.240	1.9796	0.7917	7.4777	19.50	58.40	0.1978

scope	period	N	u*	eta	delta	C	ln N	N_eff	F	eta ln N	Q_inf	tau	t_95	h
an institution	quarter	25	0.40	0.22	0.030	0.30	3.2189	10.600	2.3609	0.7082	7.3129	32.83	98.35	0.1888
a city	year	60	0.35	0.15	0.020	0.18	4.0943	21.650	3.0750	0.6142	14.062	49.50	148.28	0.1529

Every row is the same recursion. Read down the `t_95` column: the desk settles in a month of days, the team in a year of weeks, the city in a century and a half of years. **That column is the reason scales cannot be managed with a single calendar, and it is the whole content of the observation that institutions are slow.** They are not slow because of their people. They are slow because `t_95 = ln(0.05)/ln(1−delta)` and their `delta` is small — which is also the only reason they persist.

Read down the h column and a second fact appears: **headroom is roughly comparable across scopes even though `Q_inf` differs by a factor of six.** Headroom is a *rate*, and rates are what a scope can afford per period; stocks are what it has accumulated. Confusing the two is Chapter 6's central warning and it recurs at every scale.

60.7 Part I — the circuit, and what it settled

Chapters 1-6 established that the recursion is an accumulator rather than a regress, that `delta > 0` makes it a contraction and therefore guarantees arrival, and that the identity with a first-order low-pass filter is exact rather than metaphorical. From that came the pole, the time constant, the settling times, the DC gain `eta/delta`, and the frequency response that makes stored order blind to fast fluctuation and sensitive to slow trend.

What Part I contributed to everything after it is the single gauge: `h = eta ln N − delta Q − C`. Below zero, no arrangement of people saves the design; at zero, hold; above zero, grow by at most `h`. **Every chapter of this manual that recommends anything checks it against `h` first**, and the reason is Corollary 3.4: if the cost term exceeds `eta ln N − delta Q`, no interaction geometry whatsoever prevents order from decreasing. That is the most important sociological statement in the book and it is a theorem.

60.8 Part II — the geometry of difference

Chapters 7-11 established that difference is counted before it is valued, and that only what is different is counted at all. `N_eff = 1 + (N−1)u` audits a headcount for redundancy; `F = ln N_eff` is the log of what survives the audit. `F = 0` for parallel participants — **the wedding of like with like is a non-event, and that is Theorem 4.3 rather than an opinion.**

Two consequences ran through everything after. First, **the ceiling `ln N` is hard**, so a period's input is bounded no matter how the period is run — which is why Chapter 2's cadence result works and why Chapter 59's calendar arithmetic bites. Second, **concavity means the second**

genuinely different framing is the largest gain available and the fortieth is glazing bills, which set the entire argument of Chapter 39 and reappears in §60.18 below as the formal content of the practice of adopting a second perspective deliberately.

60.9 Part III — engineering space

Chapters 12-17 took the room as a circuit: `delta` for a room is how fast its arrangement reverts, `eta` is how much of what happens in it becomes furniture, `C` is what it costs to keep. Sabine's `RT60 = 0.161 V/A` is real acoustics and enters as a constraint on whether the room can carry interaction at all. Hall's proxemics sets the distances at which particular kinds of interaction are possible. Alexander's semilattice — "a city is not a tree" — is the observation that overlapping structure beats strict hierarchy, which in this notation is the observation that a tree destroys u by construction.

What Part III contributed is that `eta` is often the cheapest parameter to move, and that it is usually moved by furniture rather than by exhortation. A room that makes the good arrangement the default converts more interaction into structure than a room that requires everyone to remember.

60.10 Part IV — the sociology

Chapters 18-23 supplied the empirical layer that the formalism does not contain. Woolley and colleagues (*Science*, 2010) found a collective intelligence factor not predicted by members' average or maximum ability but correlated with social sensitivity and **equality of conversational turn-taking** — which is an `N_eff` measurement by another route. Centola and colleagues (*Science*, 2018) measured a committed-minority tipping point near 25%. Complex contagion needs reinforcing contacts and clustered topology. Ostrom's eight design principles are, read here, a specification for keeping institutional `delta` low without raising `C` past the break-even.

What Part IV contributed is the warning that the formalism is not self-sufficient. `u`, `eta` and `delta` are empirical quantities in real groups, and the literature that measures them is a science with a replication problem, not a source of settled constants.

60.11 Part V — the tradecraft

Chapters 24-28 are the procedure: survey the system, preregister the prediction, intervene, and read the failure catalogue when it does not work. **The load-bearing element is the shuffled-input control**, which appears in every falsification suite in this manual and is the control most often omitted elsewhere. If a model predicts a system's stored order as well from shuffled inputs as from real ones, then timing carries no information and the entire cadence apparatus does not apply to that system.

Part V is what separates this manual from a worldview. A framework that cannot fail cannot be trusted, and Chapters 25 and 55 name exactly how this one fails.

Part V also supplied the manual's characteristic move, which recurs at every scale and is worth naming as a method rather than leaving as a habit. **Whenever the manual encounters a phenomenon described in the language of morale, culture, energy, or will, it computes the arithmetic first and asks about the people second.** Chapter 2's premature review is not a motivation problem but a pole. Chapter 23's ossified standards body is not badly run but has a fallen

ceiling and an unchanged maintenance load. Chapter 59's offsite whose effects "wore off" did not wear off; it decayed at `(1-delta)^t` while `C` continued. **The arithmetic is not a substitute for judgement about people. It is the thing that has to be cleared before judgement about people is worth having**, because in every one of those cases the arithmetic alone accounted for the observation and the judgement would have been wrong.

60.12 Part VI — the essence law

Chapters 29-31 established that the published recipe collapses exactly to `f(n) = k n^2` with `k = 0.483425732988414` Hz, that every interval in the system is a ratio of squared integers so that `k` cancels from all relationships, and that an essence set indexed `1..N` yields exactly `ceil(N/2)` pitch classes because pitch class depends only on the odd part of the index. The bridge to the surplus law is exact: `N_eff = e^F` is the cross-disciplinary essence index, so `f = k e^{2F}` and **one nat of surplus is exactly 2.885390 octaves.**

The honesty clause is absolute and is restated here because this is the chapter most likely to be quoted out of context: the output is an essence frequency, not a physical resonance. Nothing — no material, molecule, element, room, or person — is claimed to resonate at any value the essence law produces. `k` is a chosen constant. The square law is a postulate. That `E_n prop to n^2` is also the eigenvalue spacing of a one-dimensional infinite square well is a real and interesting coincidence of form and is **not** evidence of a physical mechanism.

60.13 Part VII — the perceptual lens

Chapters 34-39 established the register everything else is written in. Weber's constant just-noticeable ratio and Fechner's `S = k ln(I/I_0)` are nineteenth-century findings replicated continuously since. Stevens' magnitude estimation produces genuine ratio-scale measurements of subjective quantity, which is the existence proof the whole volume needs: **subjective magnitudes can be measured, and the exponent a in `psi = k I^a` is a measured conversion key between registers.**

From that comes the Lens Theorem: **if perception is logarithmic, addition in perception is multiplication in the world.** Equal perceptual steps are equal world ratios, never equal world differences. A "moderate" perceived improvement is an order of magnitude produced. And Paper A's strict concavity has the same shape as Fechner's law, so **diminishing returns is not pessimism; it is what a logarithm looks like from inside.**

The moment equation `M = (D/P)E` supplies the diagnostic. `r = D/P` with distortion `|ln r|`: over-framed above one, under-framed below, and **the fixed point `D = P` is the formula's most important content** — frame matching intake, moment equalling environment, veridical contact stated as an equation.

60.14 Parts VIII and IX — the two formularies

Chapters 40-51 are the lookup layer, and the join between the two master formulas is stated there and is worth restating: **the moment feeds the storage law as its input.**

```
 M_t = (D_t / P_t) E_t                          the moment, as lived
Q_{t+1} = (1 - delta) Q_t + eta M_t - C_t     the moment becoming character
```

Perspective and perception produce the moment; the moment becomes character. From the same corpus: `L = max(0, M − C)`, the launch budget — **the dream of elsewhere is manufactured only from surplus, which is why "get out of your comfort zone" reverses the order of operations**; and `Delta_R = |ln(M'/M)|`, the return difference, in which zero means no information came back at all.

At interpersonal scope the same rows read differently and the most useful is `D/P` as **doctrine-to-data**: an over-framed institution sees what its model predicts; an under-framed one drowns in data it cannot read. Both are measurable and both are common.

60.15 Part X — the validation program

Chapters 52-56 are the manual's answer to the question of whether any of this is science. The answer given there is precise and unflattering to both sides of the usual argument. **Psychology and sociology are not pseudoscience; they are sciences with a measurement problem and a replication problem.** The Open Science Collaboration (*Science*, 2015) found 97 of 100 original studies reporting significant results against roughly 36% of replications, with effect sizes roughly halved.

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does. The Formula Validation Protocol's seven gates are the route, and FVP-7 — the cross-individual recovery test — is the gate that decides the question and the one most often skipped. The Rasch model's specific objectivity, in which a person's measure does not depend on which items were used and an item's measure does not depend on who answered, is the closest psychometrics comes to physical measurement and is exactly the property this volume requires.

Part X is also where the volume records what it has not achieved, and that record is short and specific. Not one formula in this manual has passed FVP-7 in the open literature, because the cross-individual recovery studies have not been run. The worked examples throughout demonstrate that the parameters *can* be estimated in a single system; they demonstrate nothing about whether the same operationalisation recovers the same parameters in a different system, with different people, in a different laboratory. **Until that is done, every quantitative claim in Volumes One and Two is a worked demonstration of a method and not a measured result about the world**, and Chapter 56 exists because saying so without also saying what it would take to fix it would be an evasion.

60.16 Part XI — the old notation

Chapters 57-59 read three technologies formally, and each reading is descriptive rather than interpretive:

- **A correspondence table is a finite relational graph.** The Tree of Life is 10 nodes and 22 paths. The tarot is a structured 78-element set. Saying so is a description of a structure, not a claim about it.
- **Gematria is a hash function**, a map from strings to integers. That is a literally accurate technical description and carries no claim that the map is meaningful. §60.20 computes what follows.
- **The sigil is lossy compression plus attention-binding**, and its geometry has an exact arithmetic invariant: a star polygon `{N/k}` decomposes into exactly `gcd(N,k)` closed strokes

of length N/gcd(N,k). Verified: {6/2} gives 2 strokes of length 3 — **the hexagram is literally two separate triangles**; {7/2} and {7/3} and {8/3} and {12/5} give one continuous stroke; {12/3} gives three squares; {9/3} gives three triangles; {24/9} gives three octagons.
- **The memory palace is the storage law with space as the retention medium**, and its advantage factorises into surplus, conversion and retention, of which retention is the largest.
- **Ritual is periodic forcing of a low-pass filter**, and cadence beats intensity for two independent reasons: the per-period ceiling, and the concavity of f.

60.17 The three handles

Reduce all of it and three handles remain. **Every operation this manual has described moves one or more of them, and the manual has found no operation that moves something else.**

```
HANDLE 1 -- THE LENS.        r = D/P, distortion |ln r|, moment M = (D/P)E
                          What you can perceive, and how much of the frame
                              you are supplying yourself.

HANDLE 2 -- THE STORE.       Q, eta, delta, C, h
                        What persists, how fast experience converts into it,
                       how fast it leaks, what it costs, what you can afford.

HANDLE 3 -- THE COMPOSITION. N, u, N_eff, F = ln(1 + (N-1)u), ceiling ln N
                         How many genuinely different things are in contact,
                        and the hard bound on what their contact can generate.
```

The three are not independent — the lens determines what N you can distinguish at all, which caps F, which drives Q — but they are **separately actionable**, which is the property an engineering handle needs.

And the maximum possible moment is `ln N`. That was not chosen; it follows from putting M on the same normalised scale as S. **The richest available moment is bounded by the logarithm of how many kinds of thing the perceiver can distinguish, which is Ashby's law of requisite variety appearing in first-person experience.** Only variety absorbs variety, including the variety of one's own life.

60.18 The audit

Eleven operations from the old notation, each mapped to a handle and a parameter, each priced where pricing is possible. **This is the audit that §60.2 requires, and it is presented as an audit of a list, not a proof of completeness.**

Operation	Handle	Parameter moved	What the formalism says it does
purification / banishing before work	2	C down, eta up	removes competing encodings; less interference per period

Operation	Handle	Parameter moved	What the formalism says it does
consecration of a place	2	`delta` down, `u_cross` up	a place with a fixed assigned meaning is a durable index
the correspondence table	3	declares `N`	fixes the per-interaction ceiling at `ln N`
gematria	3	a hash into `R` buckets	creates adjacency by collision; §60.20
the sigil	2, 3	`eta` up; `gcd(N,k)` strokes	compression to one retrievable token; exact stroke count
invocation / assuming a second framing	3	`k` -> `k+1` framings	`F_int = ln(1 + (k−1)u_bar)`; the second is the largest gain
the memory palace	2	`u` up, `eta` up, `delta` down	Chapter 59; the product is the gain
the liturgical calendar	2	periodic forcing	Chapter 59; cadence, ripple, and the ceiling
initiation / degree	2	`C` up, `delta` down	costly-signal break-even, Chapter 59 §59.15
the talisman / physical token	2	`delta` down	external storage; see below
divination by random draw	3	forces a non-parallel framing	expected surplus computed exactly, below

Four of the eleven can be priced exactly, and are.

The correspondence table declares `N`, and nothing else about it changes the ceiling:

table	`N`	`ln N` (nats)	bits	octaves
the four elements	4	1.3863	**2.0000**	**4.0000**
the classical planets	7	1.9459	2.8074	5.6147
the Tree of Life, nodes	10	2.3026	3.3219	6.6439
the zodiac	12	2.4849	3.5850	7.1699
the Tree of Life, paths	22	3.0910	4.4594	8.9189
nodes and paths together	32	3.4657	**5.0000**	**10.0000**
the I Ching hexagrams	64	4.1589	**6.0000**	**12.0000**
the tarot, full deck	78	4.3567	6.2854	12.5708

The bold entries are exact because those cardinalities are powers of two: **the 64 hexagrams are exactly six bits, which is not a mystery about the I Ching but a restatement of the fact that**

a hexagram is six binary lines. The framework's contribution is only to note that six bits is also `4.1589` nats and `12.0000` octaves, and that these are the same number in three registers.

The talisman is a `delta` reduction and nothing else. At a fixed conversion `eta S = 0.30` per week:

storage medium	`eta S`	`delta/wk`	`Q_inf`	`t_half`	vs held in mind
held in mind only	0.3000	0.060	5.0000	0.22 yr	1.00
written once, filed	0.2100	0.010	21.0000	1.33 yr	4.20
carried, seen daily	0.2550	0.004	63.7500	3.33 yr	12.75
engraved, in the room	0.2400	0.002	120.0000	6.66 yr	24.00

The object stores; it does not act. Its entire advantage sits in one column, and the column is `delta`. Note that the written form has the *lowest* `eta` in the table — a filed document converts less of each occasion than an active memory does — and still wins by a factor of four, because retention is a gain term. This is Chapter 3's rule and Chapter 23's rule and Chapter 59's rule, arriving from a fourth direction.

Invocation of a second framing has a computable value. `F_int = ln(1 + (k−1)u_bar)` over k framings held by one observer with mean pairwise non-parallelism `u_bar`:

k framings	`F_int` at `u_bar = 0.7`	marginal gain	% of the k=12 value
1	0.0000	—	0.0%
2	0.5306	**+0.5306**	24.5%
3	0.8755	+0.3448	40.5%
4	1.1314	+0.2559	52.3%
6	1.5041	+0.1691	69.5%
8	1.7750	+0.1263	82.0%
12	2.1633	+0.0839	100.0%

The second framing delivers 24.5% of what eleven more deliver. Deliberately adopting one other genuine perspective is the highest-return operation available to a single mind, and the twelfth is worth 15.8% of the second. That is the formal content of every tradition's insistence that the practitioner learn to see from somewhere other than where they are standing — and of Chapter 39's warning that the fortieth framing is a glazing bill.

A random draw from a correspondence table has an exact expected value. If the drawn framing's non-parallelism with the operator's current one is uniform on `[0,1]`, then

```
E[F] = integral_0^1 ln(1 + (N−1)u) du = (N ln N − N + 1)/(N − 1)
```

N	E[F]	ceiling ln N	as % of ceiling	E[N_eff] = e^{E[F]}
4	0.8484	1.3863	61.2%	2.3359
10	1.5584	2.3026	67.7%	4.7513
22	2.2382	3.0910	72.4%	9.3768
64	3.2249	4.1589	77.5%	25.1510
78	3.4133	4.3567	78.3%	30.3650

A randomiser drawing from a 78-element table returns, on average, 78.3% of the ceiling. That is what a divinatory device is worth in this formalism, stated exactly and with no claim whatsoever about the source of the randomness. **It is the same function Chapter 11 assigns to a red team and Chapter 9 to a deliberately recruited dissenter: a generator of framings the operator would not have produced.** The randomness may come from cards, from dice, from a table of numbers, or from a colleague who disagrees; the formalism does not distinguish among them and does not need to.

60.19 What the audit does not prove

Eleven operations, three handles, no residue — and that is a finding about eleven operations.

It is not a proof that no twelfth operation exists which requires a fourth handle. The audit's honest form is: *of the operations this manual examined, each mapped to a named parameter, and each mapping produced a direction and, in four cases, a magnitude.* **A reader who wishes to refute the thesis of §60.1 should exhibit an operation that does work, is agreed to do work, and maps to nothing.** That is a clean falsification target and the manual states it deliberately.

Two candidate residues are worth naming, because they are the ones a careful reader will raise and because neither survives inspection as a *fourth handle*:

Candidate one: the felt quality of the operation. Practitioners report that the operations feel like something, and the formalism does not model that feeling. **True, and out of scope rather than unaccounted for.** The moment equation gives `M` as a magnitude, not a quality; Chapter 37 states the methodology plainly — you postulate a structure, derive consequences, and test the consequences, because the inner state is not directly measurable. **That is a limitation of the instrument and it is declared, not hidden.**

Candidate two: the effect on the world outside the practitioner. This is the one the thesis refuses. Every mechanism in the audit runs through a human being who perceives, encodes, retrieves, or acts. **No operation in the table has a path to any outcome that does not pass through a person, and the manual claims none.** A practitioner whose attention, memory, framing, and cadence have all been engineered will act differently and will therefore produce different outcomes — which is a completely ordinary causal chain and is the only one on offer here.

Candidate three: the coherence of the system as a whole. A practitioner may object that the operations in §60.18 are not independent moves on separate parameters but a single integrated practice, and that the audit destroys the object by decomposing it. **This objection is the most serious of the three and the formalism has an answer that is neither dismissal nor concession.**

Paper A, Theorem 3.4 gives subadditivity on the parameter u: the surplus of a combination is at most the sum of the parts' surpluses and is generally strictly less. **Formally, then, an integrated practice is *less* than the sum of its decomposed operations, not more** — Chapter 59 §59.6 computed exactly this for the memory palace, where the additive reading of place and image overstated the truth by up to 2.25 nats. So the objection cannot be that decomposition loses a surplus, because the mathematics runs the other way.

What decomposition genuinely does lose is **the cost structure**. Eleven separate operations carry eleven per-occasion costs `c_0`; one integrated practice carries one. Chapter 59 §59.12 showed that the per-occasion fixed cost is what determines the optimal cadence and can determine whether a practice is affordable at all. **An integrated practice is cheaper to run than its decomposition, and that is a real advantage, and it is a `C` advantage rather than an `S` advantage.** The audit is therefore correct as an audit and incomplete as a design: **it tells you what each operation does; it does not tell you what it costs to do them separately, and the answer to that is usually "more."**

60.20 The restraint, worked: what gematria's correspondences are worth

This is the chapter's demonstration that restraint is a sharper instrument than permission, so it is computed rather than asserted.

Gematria maps a string to an integer. As an operation this is exactly a hash, and hashes have a governing arithmetic. For `M` strings mapping into `R` reachable values with probabilities `p_i`, the expected number of colliding pairs is `C(M,2) * sum_i p_i^2`, and **`sum p_i^2` is minimised when the distribution is uniform**, so `C(M,2)/R` is a *floor*. Real letter-value distributions are far from uniform, which makes collisions strictly more common than the floor:

corpus size `M`	reachable values `R`	floor on expected colliding pairs	P(a given word has at least one twin)
500	400	311.9	0.713226
2,000	1,000	1,999.0	0.864665
8,000	1,000	31,996.0	0.999666
8,000	1,500	21,330.7	0.995177
30,000	1,500	299,990.0	1.000000

And the birthday point — how many words you need to draw before a coincidence is more likely than not:

reachable values `R`	exact `M` for P(match) > 0.5	approximation `1.1774 sqrt(R)`
100	13	11.77
400	24	23.55
1,000	38	37.23
1,500	46	45.60
10,000	119	117.74

Draw thirty-eight words from a scheme with a thousand reachable values and a numerical coincidence is more likely than its absence. In a corpus of eight thousand words, essentially every word has a partner — the probability is 0.9997.

Therefore the discovery that two words share a value is not evidence of anything, and the amount of non-evidence is computable in advance. This is the strongest negative result in Part XI and it is delivered by arithmetic rather than by disdain.

And it does not abolish the operation. A hash that assigns a word to a bucket is a real encoding device: it produces an addressable index, and Chapter 59 showed exactly what an addressable index is worth. **What the arithmetic removes is the *evidential* reading, not the *mnemonic* one.** The operation survives as a technique for binding items to a retrieval structure and dies as a method of discovering hidden relations, and the framework says precisely which of the two it is being used as by asking one question: *is the practitioner claiming the collision was informative, or using it to remember?*

60.21 Why the restraint is the source of the power

The four refusals of CANON §5 look like limits and function as instruments.

A claim that cannot fail cannot inform. If the thesis of §60.1 permitted supernatural efficacy, no audit could test it, because any residue could be assigned to the permitted category and the audit would never close. **The refusal is what makes §60.19's falsification target exist at all.**

A framework that explains everything explains nothing. Chapter 55's insistence on the shuffled-input control, Chapter 25's preregistration, and Chapter 52's seven gates exist so that this apparatus can be wrong in a specified way. **Every one of those gates is a place where the manual has arranged to lose.**

And the strongest result in this volume is a negative one. §60.20's collision arithmetic tells a practitioner more that is useful than any permissive reading could, because it converts a vague suspicion into a number they can compute before they draw the thirty-eighth word.

The mechanism is sufficient. Nothing needs to be added to it, and the reason for saying so is not that additions are forbidden but that in eleven operations, none was required.

The Specification

S60.1 — Publish the ledger. Any deployment of this framework shall reproduce §60.3's three-way sort — proved, postulated, requires measurement — for its own claims, and shall not present a postulate in the same register as a theorem.

S60.2 — Badge in the same sentence. Per CANON §2, a postulate shall carry its badge in the sentence that states it, not in a footnote, an appendix, or a preface.

S60.3 — Name the handle. Every proposed intervention shall name which of the three handles it moves and which parameter, and shall state the expected direction. An intervention that cannot name its parameter is not specified and shall not be approved.

S60.4 — Compute before adopting. Every intervention shall be checked against `h = eta ln N − delta Q − C` before adoption. Where `h < 0` no intervention based on interaction geometry shall be approved, and the design record shall say so in those words.

S60.5 — Declare the period and the scope. Every parameter shall carry its period (S2.1) and its scope — internal or external, and which row of the master table it occupies.

S60.6 — Preregister and freeze. Every claim of effect shall be preregistered per Chapter 25, fit on a training window, frozen, and reported on held-out data against four baselines, of which one is the **shuffled-input control**. Omission of the shuffled-input control invalidates the claim.

S60.7 — Compute the chance rate before reading a coincidence. Where a correspondence, collision, match, or numerical agreement is offered as evidence, the expected number of such agreements under the null shall be computed and reported alongside it (§60.20). An observed rate not exceeding the chance rate shall be reported as no evidence.

S60.8 — Distinguish mnemonic use from evidential use. Any operation used as an encoding device shall be recorded as such, and shall not be cited elsewhere as a discovery procedure.

S60.9 — No path that does not pass through a person. Every causal claim shall name the human perception, encoding, retrieval, or action through which it operates. A claim with no such step shall be refused, per CANON §5.1.

S60.10 — Preserve the ability to decline. No intervention shall be designed, recommended, or deployed that removes a participant's ability to decline it. This applies to cadence, to requirements, to thresholds, and to every operation in §60.18's audit, and it is not subject to trade-off against any quantity in this manual.

S60.11 — State the falsification target for the synthesis. Any restatement of §60.1's thesis shall carry §60.19's refutation condition — an operation that does work and maps to nothing — in the same document.

S60.12 — Report the audit as an audit. No claim of completeness shall be made for §60.18. The number of operations examined shall be stated whenever the result is cited.

The Failure Modes

THE NOTATION FALLACY — concluding that because the old notation and this one describe the same relations, the old notation was secretly doing modern mathematics, or that the modern notation inherits the old one's authority.

- *Parameter damaged:* none. The error is in the direction of inference.
- *Diagnostic signature:* the phrase "the ancients knew" appears, or a formal result is offered as vindication of a tradition, or a tradition is offered as corroboration of a formal result.
- *Measurement:* ask which of the two is being used to license a claim about the other. A correspondence licenses claims about neither.
- *Fix:* state the correspondence as a correspondence. §60.1's thesis is about *this* manual's description, not about what anyone historically believed.

THE PERMITTED RESIDUE — allowing one operation to remain unmapped "for now," on the grounds that the framework is young.

- *Parameter damaged:* the falsifiability of the whole thesis.
- *Diagnostic signature:* an operation is described as "not yet formalised" across successive revisions, while the thesis of sufficiency continues to be asserted.
- *Measurement:* count unmapped operations. The correct number for a sufficiency claim is zero, and §60.19's target is exactly this.
- *Fix:* either map it or withdraw the sufficiency claim for the class it belongs to. **An unmapped operation is not a gap in the theory; it is a refutation waiting to be written up.**

THE PARAMETER AS TARGET — a measured `u`, `eta`, `delta` or `h` becomes a performance metric and ceases to measure anything.

- *Parameter damaged:* all of them, and the instrument as well.
- *Diagnostic signature:* reported `u` rises without any change in the roster or the disagreement structure; `delta` improves in every quarterly report. Goodhart's Law and Campbell's Law.
- *Measurement:* re-derive the parameter by an independent protocol and compare. If the two diverge and only the reported one improves, the metric has been captured.
- *Fix:* separate measurement from evaluation. The survey of Chapter 24 is run by people who are not accountable for the result, or it is not run.

THE SCALE CONFUSION — a result established at one scope applied at another without re-measuring the parameters.

- *Parameter damaged:* all of them, silently, because the *form* transfers and the *values* do not.
- *Diagnostic signature:* a personal practice's cadence imposed on an institution, or an institution's review interval applied to an individual. §60.6's `t_95` column spans 28 days to 148 years.
- *Measurement:* check that the period, `N`, `u`, `eta`, `delta` and `C` were all re-measured at the new scope. Hyperrecursive symmetry is a claim about model form only (CANON §1.4).
- *Fix:* re-measure. The equation transfers for free; nothing else does.

THE SUFFICIENCY OVERREACH — reading §60.1's thesis as a claim that the formalism explains experience, meaning, or value.

- *Parameter damaged:* none; the error is one of scope.
- *Diagnostic signature:* the framework is used to adjudicate what someone should want, rather than to compute what a wanted thing would cost.
- *Measurement:* ask what the claim would predict. A claim that predicts nothing measurable is not in this framework's domain.
- *Fix:* return to Chapter 37. **The formalism postulates a structure and tests consequences. It has no access to the inner state and says so, and a manual that computes the cost of a thing has said nothing whatever about whether the thing is worth its cost.**

THE CLOSED LEDGER — a deployment that stops publishing which of its claims are postulates.

- *Parameter damaged:* the reader's ability to check anything.
- *Diagnostic signature:* a summary, a slide, or a derived document in which the badges have been dropped for brevity. This is the most common failure and it usually happens in the second generation of a document rather than the first.
- *Measurement:* take any derived document and attempt to reconstruct §60.3's three-way sort from it alone. If you cannot, the ledger has closed.
- *Fix:* re-attach the badges. **CANON's house rule is that the mysticism is permitted in the open, with its badge on; a badge removed for brevity converts a declared postulate into an undeclared one, which is the one thing this manual has refused from its first page.**

Worked Example

The Calpine Studio: one operator, five colleagues, one year, every part

One system, carried through all eleven parts of the manual with real numbers. **Every parameter value below is a design value obtained by the protocols named beside it and is badged as illustrative; none is a measurement of any real population.**

[II] The roster audit — the geometry of difference.

```
   N = 8 declared domains      ln N = 2.0794 nats = 3.0000 bits = 6.0000 octaves
     u* = 0.55                          (Chapter 9 disagreement and error-
decorrelation audit)
    N_eff = 1 + 7(0.55) = 4.8500
    F = ln 4.8500 = 1.5790        (75.9% of the ceiling)
```

Six people declaring eight domains buy 4.85 effective ones; 3.15 are redundant. Two-source split (Theorem 4.1) with `u_cross = 0.40` between disciplines and `u_div = 0.15` within them:

```
    f(u_cross) = 1.3350    f(u_div) = 0.7178    f(u) = 1.5790
    bound:  1.3350  <=  1.5790  <=  2.0528
```

The additive figure 2.0528 exceeds the truth by 30.0%, and per S59.4 it is not used.

[I] The circuit — storage parameters. Period: one week.

```
    eta = 0.35 (Ch 4)   delta = 0.055/wk (Ch 3)   C = 0.13/wk (Ch 5)
    eta F      = 0.5526/wk        ceiling eta ln N = 0.7278/wk
    net drive  = 0.4226/wk        Q_inf = 0.4226/0.055 = 7.6844
    DC gain    = eta/delta = 6.3636      Thm 3.3 cap = 13.2328
    pole 0.945   tau = 17.677 wk   t_half = 12.253 wk
    t_95 = 52.96 wk    t_99 = 81.41 wk
  h at Q_inf = 0.7278 - 0.055(7.6844) - 0.13 = 0.1752/wk    h = 0 at Q = 10.8692
```

Per S2.4 the first honest review is week 53. Per S60.4, `h = 0.1752 > 0`, so growth is available and `r_max = 0.1752` per week is the largest increment that may be promised.

[III] The room. 7.0 x 5.0 x 3.0 m gives `V = 105.0` m³ and `S = 142.0` m². Sabine:

average absorption `alpha`	`A` (sabins)	`RT60 = 0.161 V / A`	verdict
0.10 bare, hard finishes	14.20	1.190 s	too live for six-way speech
0.20 furnished, some soft	28.40	**0.595 s**	in band
0.32 treated for speech	45.44	0.372 s	slightly dead
0.45 over-damped	63.90	0.265 s	conversation feels effortful

Target band 0.4-0.7 s requires `alpha` between 0.170 and 0.298. **The room is specified at `alpha = 0.20`.** Six seats at 1.6 m spacing — inside Hall's social distance of 1.2-3.6 m — require a circumference of 9.60 m and therefore a table 3.06 m across.

[IV] The sociology. Measured turn shares across four sessions: 34, 22, 17, 13, 9, 5 per cent.

```
turn entropy H     = 1.6329 nats        max ln 6 = 1.7918
evenness           = 91.1%
effective speakers = e^H = 5.1185 of 6
```

Woolley and colleagues' turn-taking correlate is therefore in reasonable shape and is **not** the binding constraint; the roster's redundancy is. Per Centola and colleagues' measured threshold, `ceil(0.25 x 6) = 2` committed members is the minimum bloc for overturning a house convention — a number the studio records so that a two-person push is recognised as a structural event rather than as a personality clash.

[VI] The essence set. Eight domains, ordered and justified per Chapter 33, indexed `n = 1..8`. `f(n) = k n^2` with `k = (81/64)/(phi+1) = 0.483425732988414` Hz — **a chosen reference constant, a postulate, badged.**

n	`f(n)` (Hz)	odd part = pitch class	j	audible (Hz)
1	0.483426	1	8	123.7570
2	1.933703	1	6	123.7570
3	4.350832	3	5	139.2266
4	7.734812	1	4	123.7570
5	12.085643	5	4	193.3703
6	17.403326	3	3	139.2266
7	23.687861	7	3	189.5029
8	30.939247	1	2	123.7570

Eight essences yield exactly `ceil(8/2) = 4` pitch classes — odd parts 1, 1, 3, 1, 5, 3, 7, 1. That is the available harmonic vocabulary and it is a structural limit, not a stylistic choice. The cross-disciplinary index is `N_eff = e^F = 4.8500`, giving `f_cross = k N_eff^2 = 11.371382` Hz, `2F/ln 2 = 4.5560` octaves above base, and 363.8842 Hz after a five-octave transposition. **Band and `j` are stated for every value, per CANON_SONIC §5. These are essence frequencies, not physical resonances, and nothing is claimed to resonate at any of them.**

[VII] The perceptual lens. `D` and `P` estimated by magnitude estimation against a common modulus (Chapter 53; FVP-2 satisfied, ratio scale with a true zero):

```
D = 62.0     P = 41.0     r = D/P = 1.5122     Distortion |ln r| = 0.4136
M = (D/P) E  = 1.5122 E
```

OVER-FRAMED by 51.2%: the studio is seeing what its model predicts. Two corrections, both computed: raise intake so that `P` rises to 62.0 — a 51.2% increase — or lower frame so that `D` falls to 41.0, a factor of 0.6613. The studio chooses the first, because per Chapter 48 an over-framed group that lowers its doctrine without raising its intake becomes an under-framed one. Target `r = 1.05` leaves a residual distortion of 0.0488, **11.8% of the original.**

Lens theorem check: a perceived +1 unit is a world factor of `e = 2.7183`. The measured distortion of 0.4136 corresponds to a world factor of 1.5122. **The correction is a ratio, budgeted as a ratio, and never promised as a linear improvement.**

[VIII / IX] The formularies. With `E` normalised to 1 and a personal maintenance cost `C = 0.9`:

```
L = max(0, M − C)
at r = 1.5122 (over-framed):  L = max(0, 0.6122) = 0.6122
at r = 1.0000 (calibrated):   L = max(0, 0.1000) = 0.1000
at r = 0.8000 (flooded):      L = max(0, −0.1000) = 0.0000
```

The under-framed state has no launch budget at all — there is no surplus to dream with. Note the uncomfortable and honest reading of the first row: **the over-framed state has the largest launch budget in the table**, which is why projection is generative and why the correction to `r = 1` costs something real. Chapter 41 treats this at length; the studio records that its calibration exercise will *reduce* the volume of new proposals and that this is the expected result rather than a failure.

Return difference on the studio's first external submission: sent `M = 1.00`, world returned `M' = 1.34`, so `Delta_R = |ln(1.34/1.00)| = 0.2927`. **Non-zero, therefore informative.** A `Delta_R` of zero would have meant the world's answer carried nothing.

[XI] The old notation, priced.

- **Correspondence table:** the studio's eight-domain scheme declares `N = 8` and therefore a ceiling of `ln 8 = 2.0794` nats per interaction. Nothing else about the scheme changes that number.
- **Sigil:** the studio's mark is drawn on eight points. `gcd(8,3) = 1` gives **one continuous stroke of length 8** — unicursal. `gcd(8,2) = 2` would give two separate squares. The mark

is specified as {8/3} because a single stroke is wanted, and that is a decision made on an arithmetic invariant.

- **Memory palace:** the 48-item corpus, encoded to 24 loci with two images each, moves Q_inf from 1.9160 to 26.9727 — **a factor of 14.08**, decomposing as surplus 1.5997 x conversion 2.0000 x retention 4.4000.
- **Calendar:** weekly session at the audited u* = 0.55 (eta F = 0.5526), monthly seminar at u = 0.62 (0.5863), annual retreat at u = 0.86 (0.6821), retreat placed in week 26 to avoid collision per S59.12.

```
  weekly + seminar in one week = 1.1390    vs ceiling 0.7278   -
> 0.4112 discarded (36.1%)
```

Three-year trajectory of the built calendar:

week	1	4	13	26	53	105	157	209
Q_t	0.4226	1.7313	4.4040	6.6679	8.1134	8.5193	8.5408	8.5419

Steady-state mean 8.5418, minimum 8.4259, maximum 8.7145, ripple 3.38%. **The calendar returns +0.8574 over the weekly session alone — an 11.16% gain** — and headroom at the steady mean is 0.7278 − 0.055(8.5418) − 0.13 = 0.1280, still positive, with h = 0 at Q = 10.8692.

[X] Validation — the frozen model and the conditions of failure.

Fit Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t on weeks 1-53 with 0 <= a <= 1 and b >= 0, freeze the coefficients, report on weeks 54-106 against persistence, direct-input, matched autoregressive, and shuffled-input baselines.

FVP gate	Status for this system
1 DIMENSIONS	D and P share a modulus, so r is dimensionless; M inherits E's units. Pass.
2 SCALE TYPE	D, P by magnitude estimation with a true zero; ratio scale, so division is permitted. Pass. u, eta, delta are ratios in [0,1].
3 LIMITING CASES	P −> 0 gives unbounded M (dogma); D −> 0 gives M = 0 (sensation without experience); E −> 0 gives M = 0. All named, all observed. Pass.
4 FIXED POINTS	D = P gives M = E; Q_inf = (eta F − c)/delta = 7.6844. Both solved and stated. Pass.
5 IDENTIFIABILITY	eta and delta are separable only from the transient; the fit window covers weeks 1-53, which is 3 tau. **Conditional pass — requires the zero-input window of Chapter 3.**

FVP gate	Status for this system
6 FALSIFICATION	The five conditions below, preregistered and frozen. Pass.
7 REPRODUCIBILITY	**NOT YET MET.** The parameters have been recovered in one studio. Cross-individual and cross-site recovery is Chapter 56's programme and has not been run. **This is stated as a gap, not glossed.**

Preregistered falsification. **This model is wrong if:**

1. the fitted a falls outside `[0,1]`, or its interval excludes `1 − delta = 0.9450`;
2. the fitted b does not exceed its pre-intervention value;
3. the frozen model does not beat the **shuffled-input control** on held-out data — in which case timing carries no information here and the entire cadence apparatus does not apply to this studio;
4. the empty-palace test recovers fewer than 20 of 24 loci in order at 90 days, which would put the `delta` attribution in the wrong subsystem;
5. `r = D/P` does not fall toward 1 after the intake intervention, which would mean over-framing in this group is not driven by the term the diagnostic assigns it to.

A week-13 reading of 4.4040 is 51.6% of steady state and is not a finding. It is recorded in advance, so that nobody has to argue about it in week 13.

What this volume has established

Psychology and sociology are one set of relations measured at two scopes. That was Volume Two's thesis, badged from its first page, and the master table of CANON_INTERNAL §5 is the argument: the same six parameters, the same recursion, the same gauge, with N as the categories one mind can distinguish or the domains a group holds, Q as character or as institution, `delta` as forgetting or as turnover, and `D/P` as frame-to-intake or as doctrine-to-data.

Four independent quantities in this framework are logarithmic — surplus, pitch, sensation, and the log form of the moment — **and therefore every conversion between them is linear in log space with exact constants.** One nat is 1.442695 bits, 2.885390 octaves, 3462.4681 cents. That is not a numerological coincidence; **the logarithm is the register perception itself operates in**, which Weber and Fechner established experimentally in the nineteenth century and which has been replicated since without interruption.

Subjective magnitude is measurable. Stevens' magnitude estimation produces ratio-scale measurements of subjective quantity that replicate, and that single fact is what licenses the entire volume: division requires a ratio scale, `D/P` requires division, and psychophysics already builds ratio scales for subjective quantities. The remaining work is not conceptual. It is the FVP's seventh gate, cross-individual recovery, and Chapter 56 costs it out.

And the old notation described the same relationships. The correspondence table declares N. The sigil compresses and has an exact stroke count `gcd(N,k)`. The palace lowers `delta`. The

calendar is periodic forcing. The second framing is worth a quarter of the next eleven combined. The randomiser returns 78.3% of the ceiling. Gematria is a hash whose collisions are computable and therefore evidentially worthless and mnemonically real. **Eleven operations, three handles, and no residue — in an audit of eleven operations, which is what it is and no more.**

If magic is the deliberate engineering of the perceptual lens, the storage of order through symbol, space and cadence, and the composition of non-parallel essences, then this manual has described its mechanism in modern notation. That is the author's thesis, it is a postulate, it wears its badge, and it claims no efficacy of any kind beyond what the storage law and the surplus law already produce — because there is nothing else in it for an efficacy to come from.

The next instrument

Volume One built the room. Volume Two turned the instrument on the person holding it. **Both volumes have been about order that a person or a group deliberately stores, and in every case the storage was chosen: a room arranged, a calendar set, a palace built, a framing adopted.**

There is a layer beneath that, and it was not chosen by anyone now living.

The essence law assigns a frequency to an ordinal position in a justified category set, and every application in Volume One's Part VI took its ordering from a discipline's own taxonomy. **But some sequences arrive already ordered, at enormous length, with an alphabet fixed before the ordering was read** — and the octave-transposition method of CANON_SONIC §5, which folds a raw essence frequency into the audible band by powers of two, was developed against exactly such a sequence, at a thirty-six-octave transposition.

Volume Three is that work: genomic audio and the biological instrument. It begins with the method this volume has been using implicitly and never named — **the concordance**, two columns set side by side, the old notation in one and the modern in the other, with the correspondence stated and its status marked. Chapter 57 built one for the correspondence table and Chapter 58 for the sigil; Volume Three formalises the register itself and then applies it to a corpus rather than to individual figures.

From there it turns to the cell — as a receiver, as a channel with a capacity, as a library — and to the question this chapter's arithmetic makes unavoidable. **The essence law assigns a frequency to an ordinal position. What happens when the ordering was not chosen by a discipline, and the alphabet is four?** What `N` means, what a defensible index is when nobody designed the sequence, what a channel capacity is when the channel is chemical, and whether any of Volume Two's internal parameters have a biological counterpart that can be measured rather than postulated — those are the open questions, and Volume Three states them as open.

It will be held to the same four refusals, and CANON_SONIC's honesty clauses will bind it more tightly than they have bound anything so far — because the temptation to claim a physical resonance is largest exactly where the material is physical, and **the one sentence that has governed every use of the essence law in this manual will govern it there too: the output is an essence frequency, not a physical resonance.**

The instrument does not change. Only `(N, u, eta, delta, C)` change, and the period.

That has been true for sixty chapters, at every scale from a desk to a city. Volume Three asks whether it is still true one layer down.